

FILED

12-28-15

9:59 PM

CEK/ek4 12/28/2015

BEFORE THE PUBLIC UTILITIES COMMISSION OF THE STATE OF CALIFORNIA

Application of San Diego Gas & Electric Company (U902M) for Review of its Safety Model Assessment Proceeding Pursuant to Decision 14-12-025.

Application 15-05-002
(Filed May 1, 2015)

And Related Matters.

Application 15-05-003
Application 15-05-004
Application 15-05-005

ADMINISTRATIVE LAW JUDGE'S RULING ENTERING STAFF WHITE PAPER AND COMBINED UTILITIES REPORT INTO THE RECORD AND SEEKING COMMENTS

Following up the Safety Model Assessment Proceeding (S-MAP) Workshop #4 on December 4, 2015, the attached reports to this Ruling are accepted into the record for this proceeding:

- 1) "Safety and Enforcement Division Staff White Paper on as Low as Reasonably Practicable (ALARP) Risk-informed Decision Framework Applied to Public Utility Safety," dated December 24, 2015 (Attachment 1); and
- 2) "Combined Utilities (San Diego Gas & Electric Company, Southern California Gas Company, Southern California Edison Company, Pacific Gas and Electric Company) S-MAP Uniformity Report Prepared for Discussion at December 4, 2015 S-MAP Workshop," dated December 1, 2015 (Attachment 2).

Comments in response to the Staff White Paper and Combined Utilities Report may be filed and served no later than Friday, January 15, 2016.

IT IS SO RULED.

Dated December 28, 2015, at San Francisco, California.

/s/ COLETTE E. KERSTEN

Colette E. Kersten
Administrative Law Judge

ATTACHMENT 1

**Safety and Enforcement Division
Staff White Paper on
As Low As Reasonably Practicable
(ALARP)
Risk-informed Decision Framework
Applied to
Public Utility Safety**

By

Steven Haine, P.E.
California Public Utilities Commission

December 24, 2015

CONTENTS

1. Introduction	3
2. History of the ALARP Principle	5
3. How Does ALARP Work?	7
4. ALARP Criterion of Gross Disproportionality	25
5. How Does ALARP Differ from Traditional Cost/Benefit Analysis?	29
6. Value of a Statistical Life	31
7. How does an ALARP Framework Relate to Risk Tolerance in Enterprise Risk Management?	35
8. ALARP Framework at the Threat and Risk Level	37
9. An ALARP Framework Among Competing Objectives and Constraints	39
10. Best Practices in an ALARP Framework	40
11. Hurdles to Acceptance of an ALARP Framework	42
12. Conclusion: Why Adopt an ALARP Framework in Rate Cases?	45
13. Recommendations and Roadmap and Required Elements of ALARP Implementation	45
14. Bibliography	48

DISCLAIMER

This report was prepared by California Public Utilities Commission (CPUC) staff. It does not necessarily represent the views of the CPUC, its Commissioners, or the State of California. The CPUC, the State of California, its employees, contractors, and subcontractors make no warrant, expressed or implied, and assume no legal liability for the information in this report. This report has not been approved or disapproved by the CPUC, nor has the CPUC passed upon the accuracy or adequacy of the information in this report.

1. INTRODUCTION

One of the intractable challenges confronting stakeholders in public utility rate cases is how to balance high expectations for public safety (and reliability) with rate affordability. Except in areas where existing safety regulations provide clear guidance based on compliance requirements, finding the right balance often devolves into subjective arguments with proposed choices that need to be adjudicated by a regulatory body. Against such a backdrop, a rigorous risk-informed decision framework able to provide an objective and analytical approach to justify the proposed safety expenditures in a rate case is needed. All stakeholders in the rate case would benefit from such a risk-informed decision framework.

This paper explores a risk management concept known as As Low As Reasonably Practicable (ALARP) and its possible application to the safety aspects of public utility operations and rate cases. In the most simplified terms, ALARP is a systematic risk-informed decision framework used to decide whether risk mitigation is needed and, when it is needed, how much should be spent until the mitigation costs are deemed to be disproportionately excessive relative to the benefits. The paper is an outgrowth of the California Public Utilities Commission's (CPUC) directive to the major California energy utilities to use risk-based decision-making in developing safety expenditures in rate cases.¹ The ALARP risk-informed decision framework represents a useful tool that could help fulfill this requirement.

Safety expenditures in traditional cost-of-service utility ratemaking rely strongly on subjective, qualitative, and sometimes arbitrary projections based on operating history and changing circumstances, supplemented by the need for compliance with safety regulations and adherence to industry safety best practices. An ALARP framework would allow a utility operator to lessen the reliance on subjective and qualitative judgment by methodically determining how much risk mitigation is needed in a way that balances safety with cost. With ALARP, regulators would benefit by being able to rely on solid, quantifiable justifications to

¹CPUC Decision 14-12-025, Decision Incorporating a Risk-Based Decision-Making Framework into the Rate Case Plan and Modifying Appendix A of Decision 07-07-004.

render a decision in rate cases that balance safety with just/reasonable rates. When optimization techniques are applied to ALARP, ratepayers would also benefit when rates would be kept as low as necessary to produce the level of safety sought by the utility operator and regulator.

The heart of an ALARP framework is cost/benefit analysis, but as will be explained in detail in this paper, ALARP is more than that. ALARP is a framework of risk management consisting of several key components that go beyond simple cost/benefit analysis. The ALARP framework implicitly assumes that risks cannot be entirely eliminated and that there is always some tradeoff between safety (as measured by residual risk after mitigation) and utility rate affordability. ALARP would help a utility operator to determine when risk from a particular threat has been sufficiently mitigated to some acceptable level, so that additional money could be spent elsewhere to more effectively reduce risks from other threats and to minimize ratepayers' total costs. When coupled with multi-attribute optimization techniques, an ALARP framework would allow a utility operator to strike an appropriate balance among safety, reliability, and utility rates.

As the adage goes, the devil is in the detail. This is definitely true of attempts to introduce new rules and new regulatory frameworks, where the danger of unforeseen effects always lurks. This paper in no way pretends to be an exhaustive treatise on ALARP. Indeed, much has been left out of the analysis. The intent of this paper is rather to provide sufficient descriptive and analytical details to spark a discussion on the benefits, drawbacks, feasibility, and hurdles of formally incorporating an ALARP risk-informed decision framework into public utility operation and rate case proceedings.

2. HISTORY OF THE ALARP PRINCIPLE

The principle of ALARP has its origin in British case law, such as *Edwards v. National Coal Board* [1949] 1 All ER 743,² which dealt with whether it was reasonably practicable to prevent any possibility of a rock fall in coal mines. The case established the concept of “reasonably practicable” and distinguished it from being “physically possible” as these terms are applied to risk mitigation measures.

In the judgment in *Edwards v. National Coal Board*, Lord Justice Asquith of the Court of Appeal stated:

“Reasonably practicable is a narrower term than ‘physically possible’ and implies that a computation must be made... in which the quantum of risk is placed in one scale and the sacrifice involved in the measures necessary for averting the risk (whether in time, trouble or money) is placed in the other and that, if it be shown that there is a great disproportion between them – the risk being insignificant in relation to the sacrifice – the person upon whom the obligation is imposed discharges the onus which is upon him.”³

The court in essence established an implied duty for any party (a duty-holder) that engages in activities giving rise to public risk to perform a risk assessment and (for candidate activities mitigating each risk) to conduct a cost/benefit analysis to compare the cost “(whether in time, trouble or money)” of risk mitigation against the benefit of the resulting risk reduction. If the cost is disproportionally high relative to the benefit, then the duty-holder is discharged from the (further) obligation to mitigate the risk, because the mitigation is then deemed not reasonably practicable, even if it be physically possible to do so. Conversely, if the cost is not disproportional to the benefit, then the duty-holder has an absolute duty to apply (and to continue to apply) risk mitigation until the cost disproportionally exceeds the benefit. This is the ALARP risk-informed decision principle in a nutshell.

² http://www.safetyphoto.co.uk/subsite/case%20e%20f%20g%20h/edwards_v_national_coal_board.htm

³ <http://www.hse.gov.uk/risk/theory/alarpglance.htm>

The ALARP principle was formally adopted in the UK by the Health and Safety at Work, etc. Act 1974 (HSWA).⁴ Outside of the UK, ALARP is not generally practiced, but several countries around the world have employed elements of the ALARP framework either wholly or partially, including Australia, Abu Dhabi, Belgium, Brazil, Denmark, Hong Kong, Ireland, Netherlands, Saudi Arabia, the U.S., and, of course, the UK.^{5, 6, 7, 8, 9}

In the U.S., the ALARP principle has been employed by the U.S. Army Corps of Engineers since 2009.^{10, 11} The U.S. Nuclear Regulatory Commission has used a principle similar to ALARP since the 1950s to regulate exposures to radiation (where it goes by the name As Low As Reasonably Achievable, ALARA). In 1972 and in 1977 the ALARA principle was adopted into two U.S. federal regulations (10CFR835 and 10CFR20) to regulate radiation exposures at Department of Energy and NRC-regulated facilities.¹² In the two federal regulations, ALARP is defined as “making every reasonable effort to maintain exposures to radiation as far below the dose limits in this part as is practical consistent with the purpose for which the licensed activity is undertaken, taking into account the state of technology, the economics of improvements in relation to state of technology, the economics of improvements in relation to benefits to the public health and

⁴ <http://www.legislation.gov.uk/ukpga/1974/37/contents>

⁵ R M Pitblado, M Bardy, P Nalpanis, P Crossthwaite, K Molazemi, M Bekaert, V Raghunathan, Det Norske Veritas, International Comparison on the Application of Societal Risk Criteria

⁶ Australian National Offshore Petroleum Safety and Environmental Management Authority (NOPSEMA) Guidance Note on ALARP: <http://www.nopsema.gov.au/assets/Guidance-notes/N-04300-GN0166-ALARP.pdf>

⁷ Danish Energy Agency's Guidelines on Safety and Health Related Conditions on Offshore Installations, etc.

⁸ Hong Kong Advisory Council on the Environment, http://www.epd.gov.hk/epd/english/boards/advisory_council/eia_paper026.html

⁹ Aven, Terje; Vinnem, Jan Erik and Vollen, Frank; “Perspectives on Risk Acceptance Criteria and Management for Offshore Applications - Application to a Development Project”; **International Journal of Materials & Structural Reliability** Vol.4, No.1, March 2006, pp 15-25.

¹⁰ Interim Tolerable Risk Guidelines for US Army Corps of Engineers Dams http://www.nfrmp.us/TRG2010/docs/USSD_2009_Corps_Tolerable_Risk_Dam_Safety_10_Feb_09.pdf

¹¹ U.S. Army Corps of Engineers, Best Practices in Dam and Levee Safety Risk Analysis - Public Risk Tolerance and Risk Guidelines.

¹² 10CFR835 “Occupational Radiation Protection” regulating ALARA in DOE facilities; and 10CFR20 regulating ALARA facilities licensed by the NRC as prescribed by 10CFR part 50, and Part 52 and Part 70 –ANS Local Section Address Eric P. Loewen 2011 that’s the commercial nuclear power plants, food irradiation facilities, medical facilities that handle radiation sources.

safety, and other societal and socioeconomic considerations, and in relation to utilization of nuclear energy and licensed materials in the public interest.”

Except for use in the nuclear industry and by the U.S. Army Corps of Engineers, the ALARP principle has yet to see wide usage in the U.S.. However, a primary component of the ALARP framework involving cost/benefit analysis to compare the incremental cost of risk reduction measures against the monetized value of a statistical human life saved by these measures has long been used by several federal government agencies, including the Department of Transportation, the Environmental Protection Agency, and the Federal Drug Administration.

3. HOW DOES ALARP WORK?

The ALARP principle conceptually divides activities¹³ into three categories based on their physical risk¹⁴ to human lives or property. At the top of the scale are activities that possess such high risk that no such activities are to be tolerated by society (except for extraordinary circumstances) unless the risk can be reduced to below a certain selected threshold. At the bottom of the scale are those activities that possess such low risk (below another selected threshold) that society generally deems the risk to be broadly acceptable. In between the top and bottom risk thresholds are those activities whose risk society will tolerate provided that the risk has been reduced by the duty-holder¹⁵ to “as low as reasonably practicable.” This generally means the duty-holder has to demonstrate that expenditures for risk mitigation have reached a point where any further spending would be considered grossly disproportionate relative to the incremental risk reduction benefits to be derived.

¹³ Under an ALARP framework, activities include both actions and all physical appurtenances directly associated with the actions, such as infrastructures, machinery, etc.

¹⁴ For the purpose of this paper, risk is a mathematical quantity defined as the expected value of frequency x consequence caused by a threat.

¹⁵ Since this paper’s emphasis is the application of the ALARP principle to public utility operation, the term “duty-holder” can simply be defined as any public utility operator. The two terms will be used interchangeably in this paper.

Symbolic representation of ALARP:

An ALARP framework is traditionally explained in risk management publications by the use of a conceptual “carrot diagram” such as one shown in Figure 1.¹⁶ In the diagram threats are segregated into three regions based on a threat’s physical risk to society:

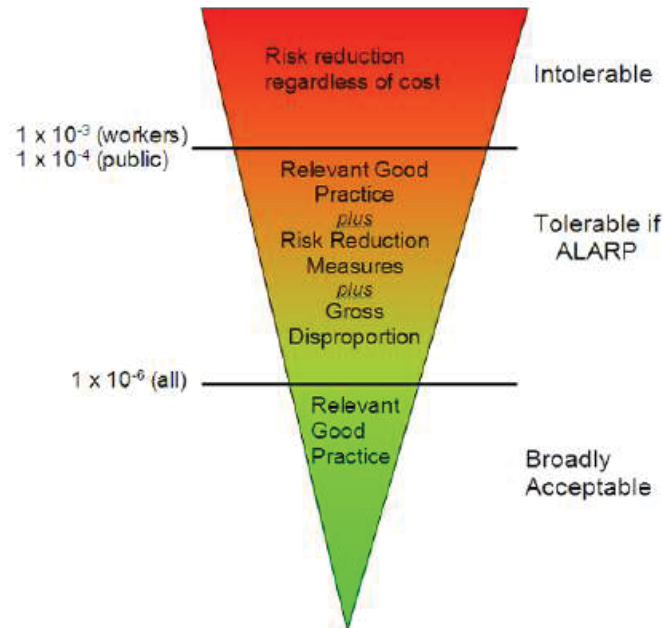


Figure 1: Symbolic ALARP diagram

The width of the “carrot” is meant to convey increasing risk along the vertical axis. The three regions are separated by two threshold limits of societal risk. The top threshold at an individual’s annual probability¹⁷ of death of 10^{-4} (1 in 10,000 chance of death per person per year, for members of the public) or 10^{-3} (1 in 1,000 chance of death person per year, for

¹⁶ HSE Guidance on ALARP Decisions in COMAH, http://www.hse.gov.uk/foi/internalops/hid_circs/permissioning/spc_perm_37/

Health and Safety Executive (HSE) is the UK government entity in charge of administering the ALARP principle. COMAH is an acronym that stands for Control of Major Accident Hazards regulations. The ALARP principle falls under the umbrella of COMAH regulations administered by the HSE.

¹⁷ An annual probability (equivalent to the concept of relative frequency in this application) of death of 10^{-4} (1 in 10,000 chance of death per person per year) is equivalent to 1 death per year per 10,000 similar individuals. It is also equivalent to 1 death per 10,000 years per individual. Likewise, an annual probability (relative frequency) of death of 10^{-3} (1 in 1,000 chance of death per person per year) is equivalent to 1 death per year per 1,000 similar individuals. It is also equivalent to 1 death per 1,000 years per individual.

employees of the utility company) represents the maximum individual accident rate that society is willing to accept for a risk from a threat, for example the potential of a fatality caused by contact with a downed energized conductor.

The threshold frequency of 10^{-4} deaths/person-year (for members of the public) or 10^{-3} deaths/person-year (for employees of the utility company) are numbers cited in the UK's Health and Safety Executive (HSE) Guidance on ALARP Decisions in the HSE's Control of Major Accident Hazard (COMAH) regulations.¹⁸ The reason for having a higher threshold rate for workers than for members of the public can be explained as follows: Workers have far more exposure to potential dangers arising from a company's operations and infrastructures than members of the public. Workers are thus naturally far more likely to experience accidents than the public simply from the higher exposure. Additionally, since the utility workers willingly take on the additional exposure to risks and since they are compensated by the company for this higher exposure, society deems the higher threshold rate for workers compared to members of the public as reasonable. This reasoning is reflected in the different threshold rates society would accept for the two groups.

Risks above the top threshold have such high risk that society deems them to be unacceptable except in extraordinary circumstances. The lower threshold line at an accident rate of 10^{-6} deaths/person-year (1 in 1,000,000 chance of death per person per year) represents the frequency limit below which society deems a threat to be broadly acceptable without any need for additional risk mitigation, except those mitigation measures needed to comport with industry best practices and regulatory requirements.

¹⁸ HSE Guidance on ALARP Decisions in COMAH,
http://www.hse.gov.uk/foi/internalops/hid_circs/permissioning/spc_perm_37/

Health and Safety Executive (HSE) is the UK government entity in charge of administering the ALARP principle. COMAH is an acronym that stands for Control of Major Accident Hazards regulations. The ALARP principle falls under the umbrella of COMAH regulations administered by the HSE.

In between these two tolerability threshold lines is the ALARP region, in which a risk is deemed acceptable only to the extent that an operator has reduced the risk to a level as low as reasonably practicable, i.e. the ALARP condition has been met for that risk. To demonstrate ALARP is satisfied, an operator has to show that any further spending on risk mitigation would be considered grossly disproportional relative to the incremental monetized value of safety reduction to be achieved.

Actual ALARP framework:

Although Figure 1 gives a good conceptual representation of ALARP, to understand how ALARP actually works, a realistic ALARP diagram, as shown in Figure 2, is needed:

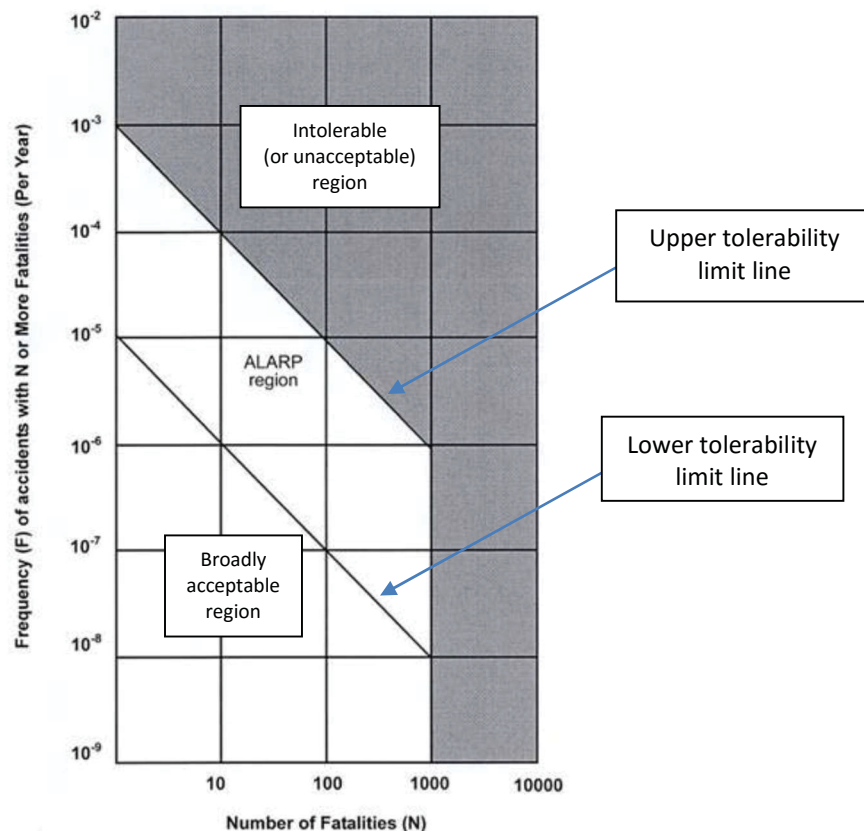


Figure 2: Realistic basic ALARP diagram

The realistic quantitative ALARP diagram in Figure 2 can also be explained by a risk matrix table concept often found in relative risk scoring models, when the risk score categories are also

based on a logarithm scale as the ALARP diagram does.¹⁹ A risk matrix table (Figure 3a), with the relative risk scores omitted from the cells, can be interpreted as a qualitative version of the actual ALARP diagram in Figure 2, once the downward sloping ALARP tolerability lines (dashed blue lines) and the ALARP region labels are overlaid onto the risk matrix table. Figure 3a shows a 4 x 4 risk matrix and the three regions (intolerable, ALARP, broadly acceptable) are demarcated by the stair step boundary lines of the cells. As the number of cells increases in a risk matrix, the boundary lines for the three regions defined by the stair steps would become finer and finer and would more closely follow the tolerability lines as in Figure 3b.

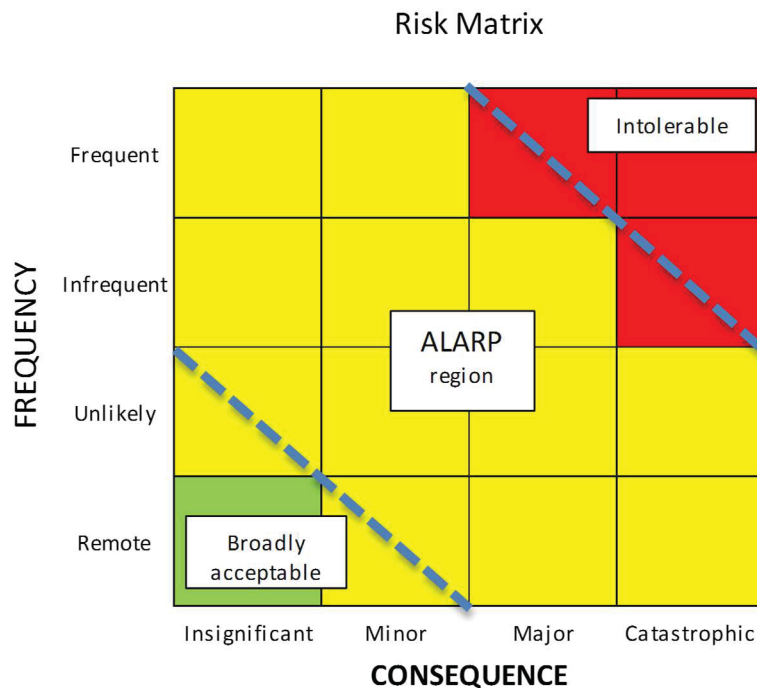


Figure 3a: 4x4 Risk Matrix with ALARP labels overlay

¹⁹ The terms relative risk ranking, scoring method, and indexing method describe the same concept of assigning relative risk scores to individual risks (or risk categories). Relative risk scores are unit-less, dimensionless values and have no physical interpretation and no meaning on their own. Relative risk scores have relative meaning only in relation to one another.

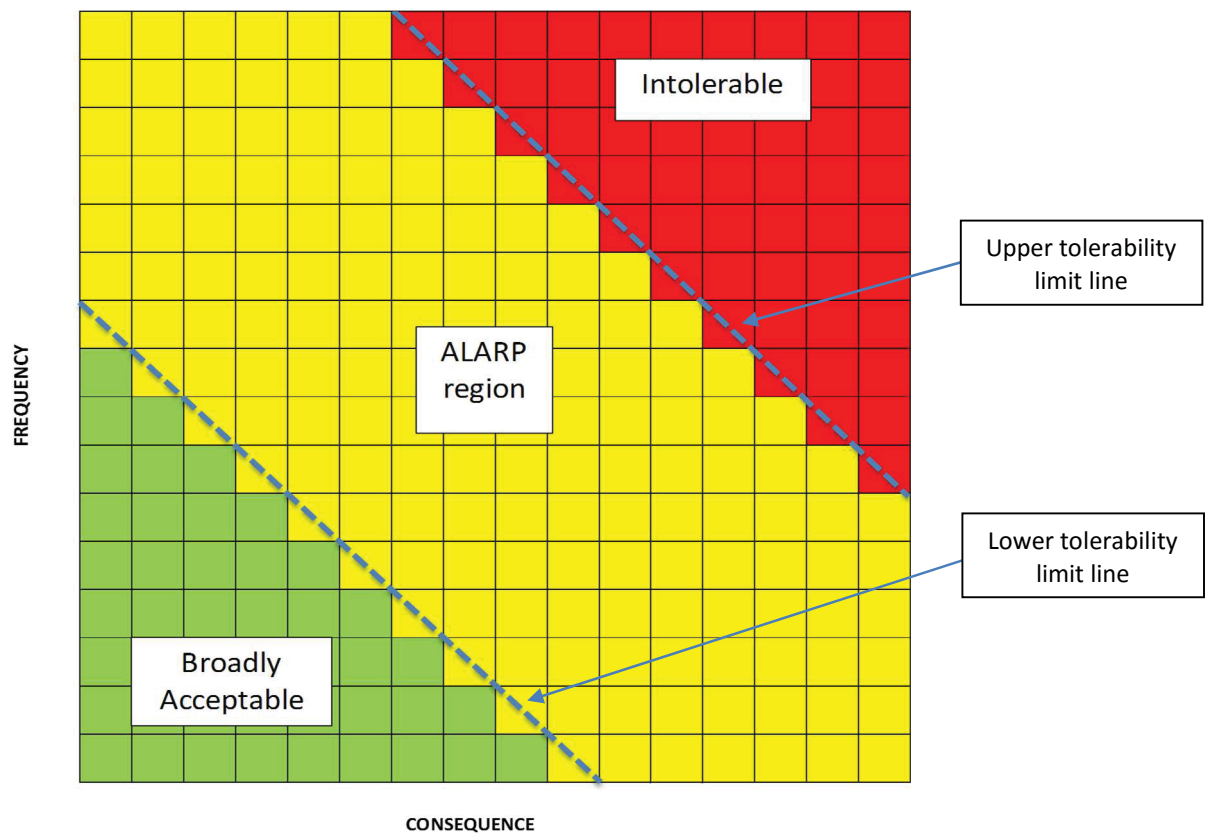


Figure 3b: 16x16 Risk Matrix with ALARP labels overlay

The primary difference between Figure 1 and Figure 2 lies in the fact that the actual threshold lines in Figure 2 are now sloping downward as a function of increasing fatalities. In Figure 2 above, there are three regions as in Figure 1, with each region representing the frequency of incidents causing N fatalities or more. The upper and lower bound frequency values of 10^{-3} deaths/person-year and 10^{-5} deaths/person-year at $N=1$ are values adopted in the UK. A regulating entity could adopt different values or in cases where an ALARP framework has not been adopted by regulation, an operator may select its own boundary values. The three regions are demarcated by boundary lines with a slope of -1 on graphing paper with a logarithm scale on both the x and y axes. The -1 slope is also a set by a regulating entity, or selected by an operator in cases where an ALARP framework has not been adopted by regulation.

The key point is that the slope of the tolerability lines must be negative. The physical interpretation of the negative slope for this boundary line is that accidents having an increasing

number of fatalities ought to be decreasingly likely in order for the risk to be socially acceptable.

Before proceeding any further we need to describe frequency-fatalities exceedance curves²⁰, or FN curves²¹. For any threat²² that can affect public safety, an FN curve plots the frequency (measured in deaths/person-year) of accidents with **N or more** fatalities per year caused by that threat on the vertical axis against different values of N on the horizontal axis. In essence, an FN curve describes the accident causing potential (measured in frequency of N or more fatalities) of an identified threat, as that threat applies to a utility operator based on the operator's unique circumstances.

A mathematical property of FN curves is that they are either flat or sloping downward as N increases but never upward sloping. A representative FN curve is shown in Figure 4. An FN curve is typically shown on graphing paper with logarithmic scales on both the vertical and horizontal axes. An FN curve can be either derived from empirical accident data or represented by probabilistic models, which in turn are usually based on empirical data. Since an FN curve represents the potential physical risk to society at a fixed point in time, an FN curve can change shape or shift based on any circumstances that affect the calculation of risk.

²⁰ In probability theory and statistics, the **cumulative distribution function (CDF)**, or just **cumulative distribution**, or **distribution function**, describes the probability that a real-valued random variable X with a given probability distribution will be found to have a value less than or equal to x . In the case of a continuous distribution, it gives the area under the probability density function from minus infinity to x . An exceedance curve is the complement of the distribution curve, i.e. $\text{Exceedance} = 1 - \text{Cumulative distribution}$.

²¹ A full discussion on FN curves can be found in this HSE publication: **Transport fatal accidents and FN-curves: 1967-2001**, <http://www.hse.gov.uk/research/rrpdf/rr073.pdf>.

²² An example of a threat in gas operation is internal corrosion.

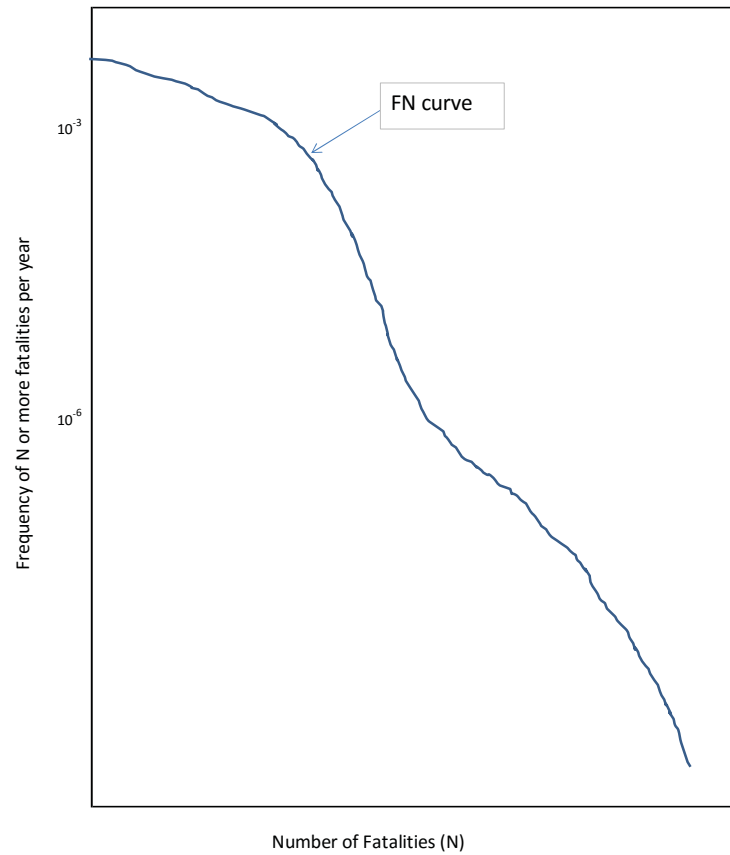


Figure 4: Representative shape of a typical FN curve

When an FN curve is constructed with the same units of measurement as an ALARP diagram, the two can be superimposed onto each other. An actual ALARP framework can now be explained in terms an FN curve overlaid on top of a realistic ALARP diagram as shown in Figure 5.

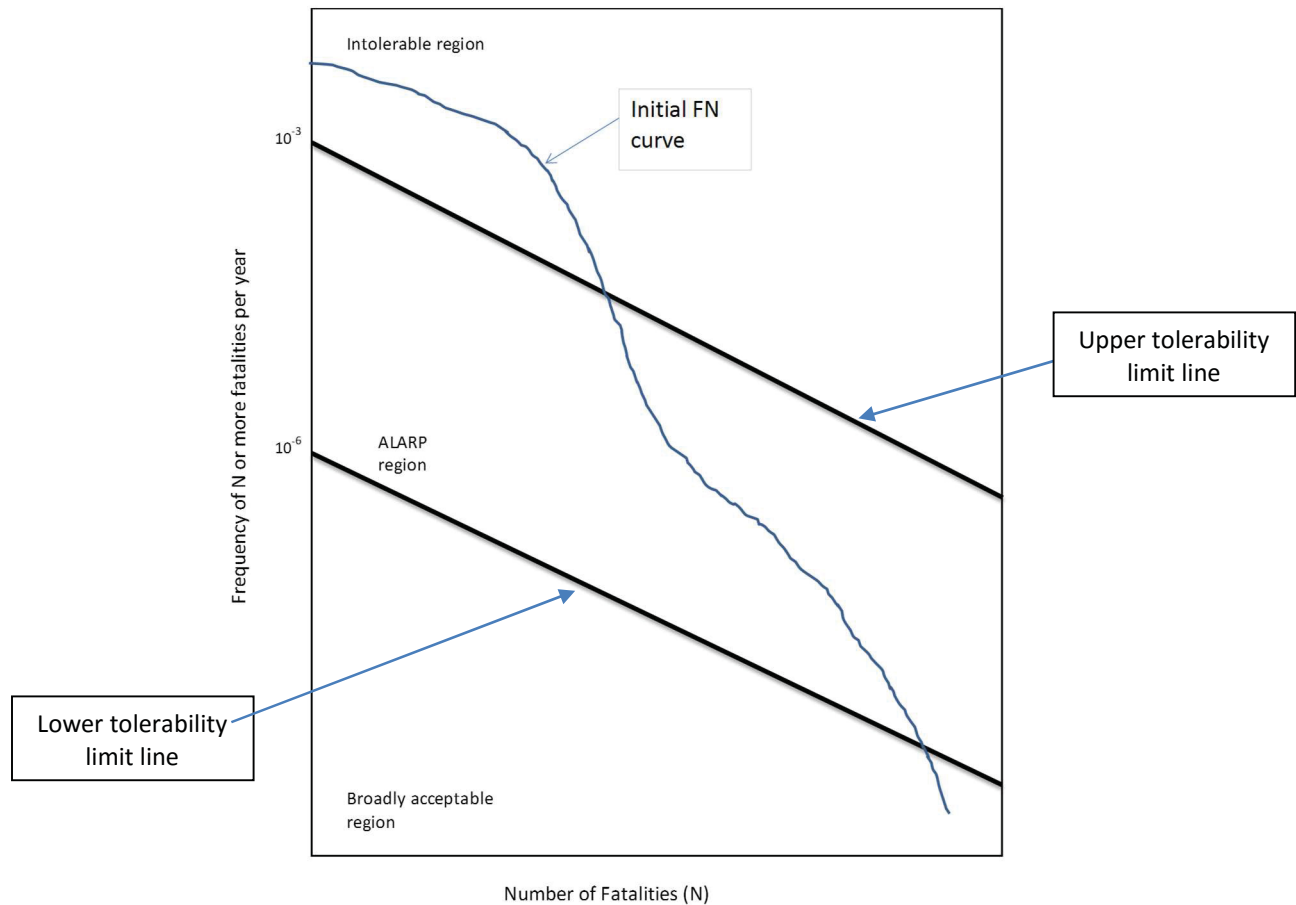


Figure 5: FN curve with ALARP diagram overlay

Unacceptable (or intolerable) region: In the unacceptable region, the unacceptability threshold frequency is 10^{-3} at $N=1$. As N increases, the unacceptability threshold begins to slope downward with a slope of -1 (representing a lower frequency and a higher safety standard). As previously stated, the physical interpretation of the negative slope for this boundary line is that accidents having an increasing number of fatalities ought to be decreasingly likely in order for the risk from the threat to be socially acceptable. Coincidentally, an FN curve also slopes down as N increases; this is a reflection of statistical data of high fatality events being typically much rarer than low fatality events.

Any threat where any part of its FN curve overlaps the unacceptable region must have the associated risk mitigated regardless of the cost of mitigation until the resulting FN curve falls

out of this region. If this cannot be achieved, then the activity giving rise to that threat must cease (or the infrastructure must be dismantled) so that the threat will no longer pose a risk to society. Referring back to Figure 5, since a portion of the pre-mitigation FN curve lies in the intolerable region, the risk is unacceptable. Risk mitigation must be applied until the FN curve falls just below the upper tolerability limit line as shown in Figure 6.

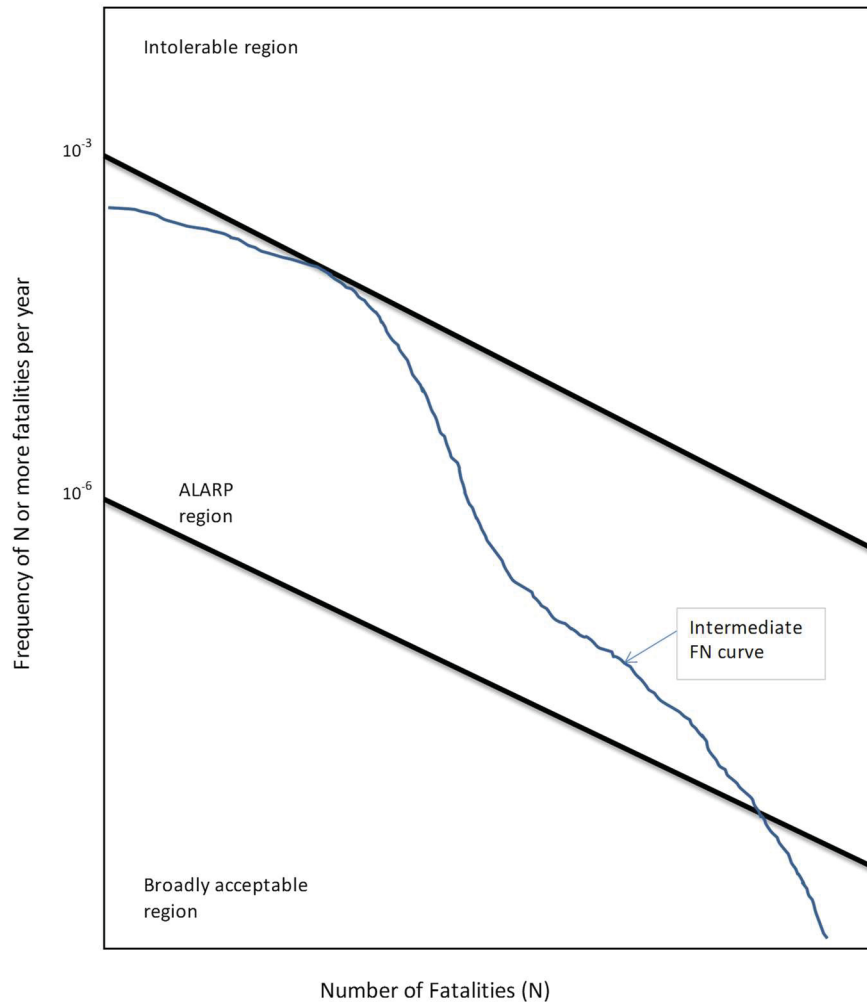


Figure 6: Just sufficient risk mitigation has been applied to move the initial FN curve in Figure 5 from the Intolerable region into an intermediate position in the ALARP region.

Finally, there may be available mitigations that might improve safety marginally, but they are so costly and have such a minor impact on risks that they really are not worthwhile pursuing. For instance, one way of reducing the risk of electrocution from a downed conductor might be to shut down the entire circuit, not just the section affected by the fault. But that might entail significant cost or inconvenience to the locality without significantly decreasing the risk.

There can be variations of this concept.²³ For example in the UK's ALARP framework, an FN curve only needs to fall below the tolerability limit line at one point (at $N=50$, $F(N)=1/5000$ per year) rather than having to fall below the entire upper tolerability limit line.

When risk mitigation measures move the resultant FN curve into a lower risk region, the rules governing the resultant region with the higher risk will apply. For example, suppose risk mitigation shifts the original FN curve from the unacceptable region into a lower position that simultaneously overlaps the ALARP region and the acceptable risk region, the rules governing the ALARP region apply since the ALARP region is the higher risk region of the two resulting regions.

Broadly acceptable region: This is the region below the boundary at a frequency of 10^{-5} at $N=1$, with a slope of -1. Threats with FN curves entirely in this region generate such low risks that society deems these threats and their associated risks to be broadly acceptable and no further risk mitigation is required as long as the FN curve remains entirely in this region.

ALARP region: In between the unacceptable region and the acceptable region is the ALARP region. Threats in the ALARP region have more potential to cause harm than those in the acceptable region, but less so than those in the unacceptable region.

²³ HSE document, *Reducing risks, protecting people* (often referred to by the acronym, R2P2), <http://www.hse.gov.uk/risk/theory/r2p2.pdf>.

For a threat with any part of its FN curve in this region, an operator must continue to reduce risks until either of the conditions apply: 1) the resulting FN curve falls entirely into the acceptable region or 2) the operator has demonstrated that expenditures on risk mitigation have reached a point at which any further spending would be considered grossly disproportionate relative to the additional risk reduction benefits to be derived from the incremental expenditure.

When condition 2 is met, the operator is described as having reduced the risk for that threat to a level as low as reasonably practicable, i.e. the ALARP criterion has been satisfied for that threat and no further risk mitigation is required (until conditions change).

In Figure 6, just sufficient mitigation has been applied to move the FN curve to just below the upper tolerability limit line, but more mitigation must be performed until the ALARP criterion of gross disproportionality is satisfied. This is illustrated in Figure 7. Figure 8 shows Figure 6 and Figure 7 combined together.

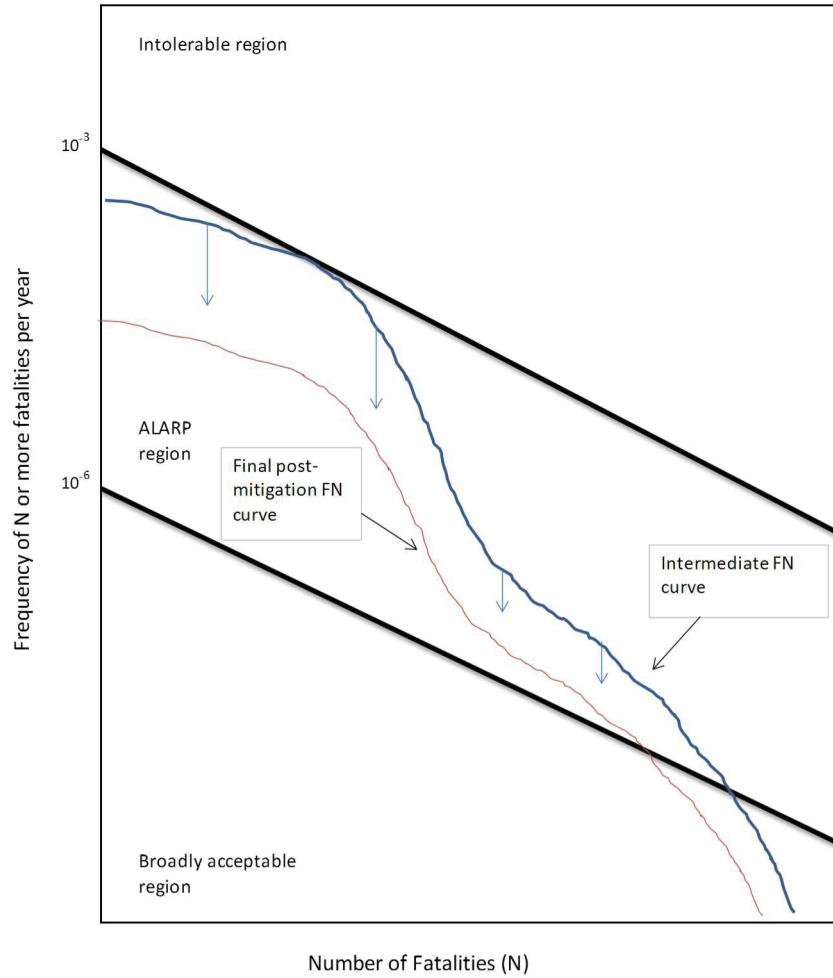


Figure 7: Additional risk mitigation applied until the ALARP criterion is satisfied. The effect is to push the FN curve down even further.

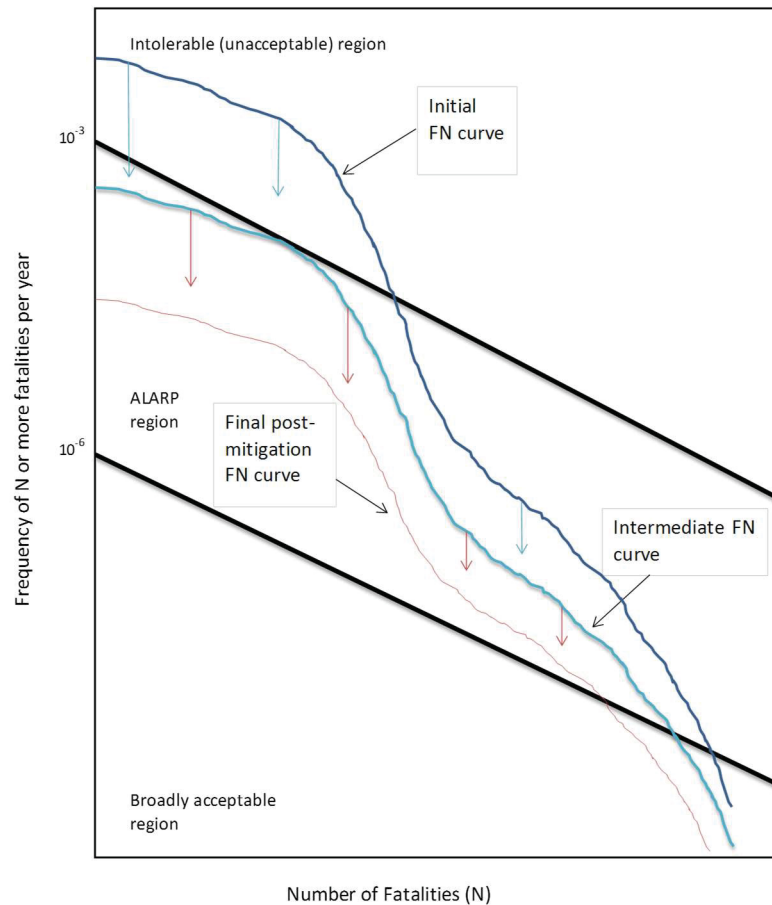


Figure 8: Graph combining Figure 6 and Figure 7

Situations may arise in which risk mitigation will push the FN curve entirely into the broadly acceptable region but the ALARP criterion is still not satisfied. In such cases the necessity to apply further risk mitigation stops when the FN curve falls entirely into the broadly acceptable region. This situation is illustrated in Figure 9:

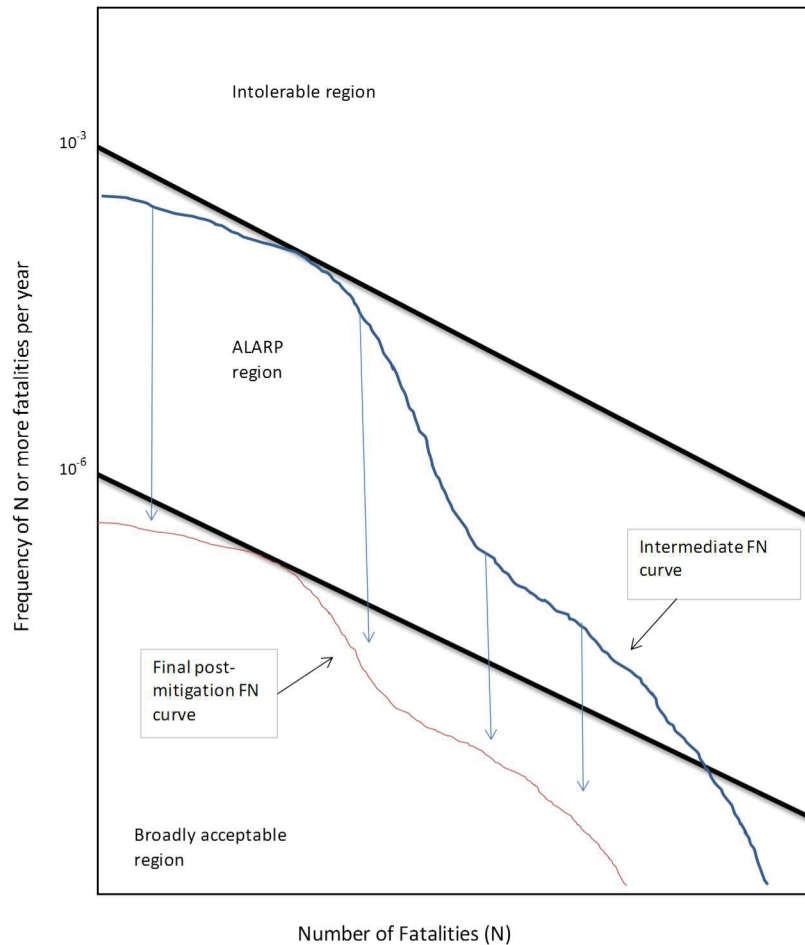


Figure 9: Case where additional mitigation pushes the FN curve to just below the lower tolerability limit line but where the ALARP criterion is still not satisfied.

To summarize what has been discussed so far:

Case 1: Any part of the initial FN curve lies in the unacceptable (intolerable) region. The unmitigated risk is considered intolerable regardless of the cost of mitigation. Risk mitigation then consists of two steps. Step 1 is to apply sufficient risk mitigation to move the initial unmitigated FN curve (Figure 5) to an intermediate position just below the upper tolerability limit line (Figure 6). Step 2 is to apply additional mitigation until either the ALARP criterion of gross disproportionality is satisfied (Figure 7 and Figure 8) or until the FN curves falls entirely into the broadly acceptable region (Figure 9), whichever comes first.

Case 2: Part of the initial FN curve lies in the ALARP region but no part of it lies in the unacceptable (intolerable) region. Risk mitigation consists of only one step. The operator must apply risk mitigation until either the ALARP criterion of gross disproportionality is satisfied (Figure 10a) or until the FN curves falls entirely into the broadly acceptable region (Figure 10b), whichever comes first. Figure 10b represents the case where the gross disproportionality ratio is still not satisfied even when the FN curve falls just below the lower tolerability limit line.

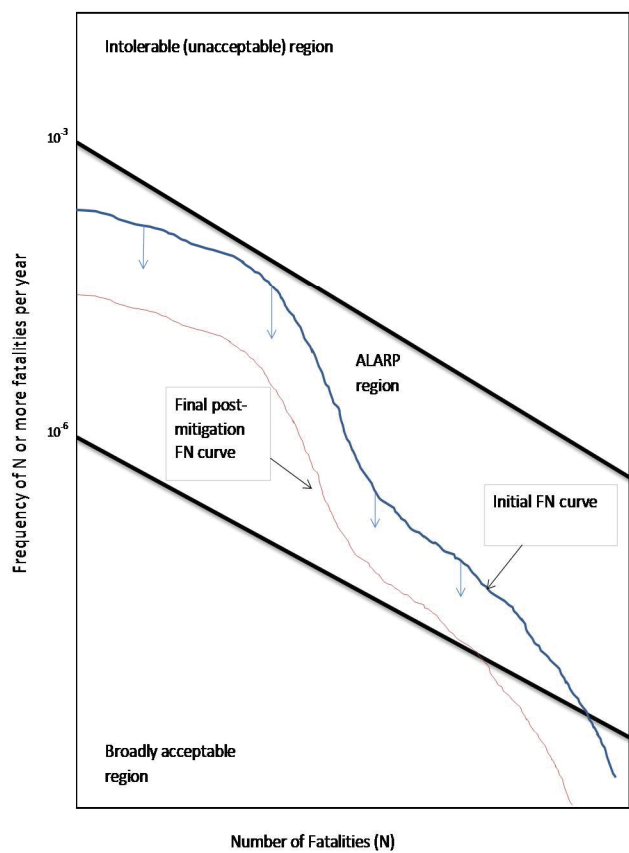


Figure 10a: Case 2 where ALARP ratio is satisfied

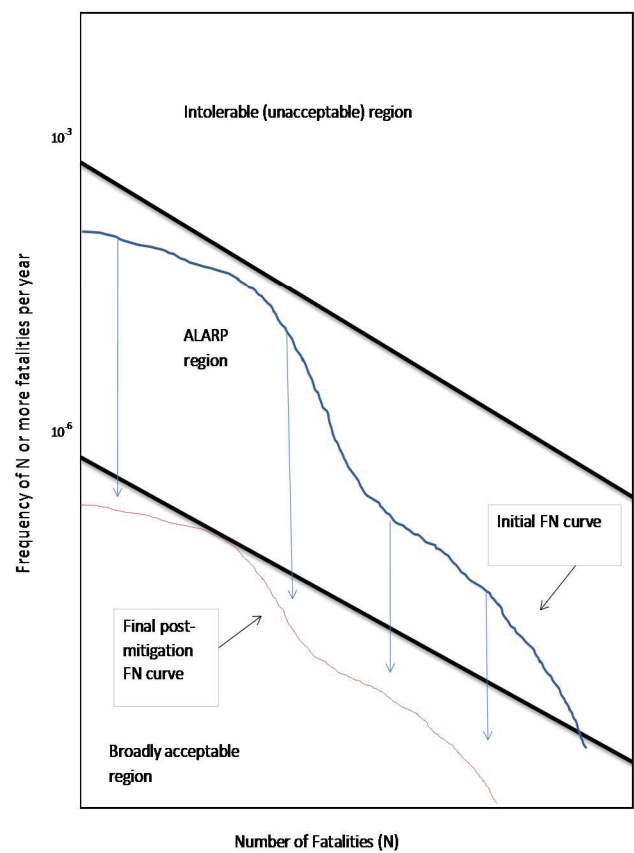


Figure 10b: Case 2 where ALARP ratio is not satisfied

Case 3: The initial FN curve lies entirely in the broadly acceptable region. No additional risk mitigation is required until conditions change.

In all cases the operator has the duty to continually evaluate the FN curves based on changing conditions that can shift an FN curve into another region and must react accordingly based on the criteria in the resultant regions.

To put the concepts developed so far into perspective, let us consider two examples:

Example 1: A system consisting of only two threats in gas transmission operation: external corrosion and internal corrosion.

The extent of external corrosion can be detected by a method known as external corrosion direct assessment (ECDA). Internal corrosion can be detected by using inline inspection tools (ILI, also commonly referred to in the gas industry as “pigging”). Under ALARP both external corrosion and internal corrosion would be subject separately to the same tolerability limit lines and the same disproportionality ratio. However, the respective methods used to mitigate the two different risks will have different costs and different levels of risk reduction benefit relative to cost. A separate cost/benefit ratio would be calculated for external corrosion and internal corrosion. ECDA and ILI activities, addressing external corrosion and internal corrosion respectively, would continue either until the ALARP disproportionality ratio is reached (Fig. 10a) or until the resulting FN curve for each threat falls entirely into the broadly acceptable region (Fig. 10b), whichever comes first.

Example 2: A gas distribution pipeline operator did not keep records of the resin specification or manufacturing date on polyethylene gas pipelines that the operator installed through the years. Several decades later, the operator became aware that many of the pipelines it installed are made of Aldyl A and that some of the Aldyl A pipelines of certain vintages and certain resin formulations are particularly susceptible to failure due to slow crack growth. Unfortunately, the state of records is such that the operator cannot retroactively ascertain the resin type or manufacturing date that are now in service. Should the operator remove all Aldyl A pipes in service from its system to reduce the risk?

Let us try to answer the question by applying the ALARP concept. The first step in applying the ALARP concept is to determine which region (intolerable, ALARP, broadly acceptable) the FN curves fall in. Since the FN curves are threat-specific, it is necessary to first identify all threats and perform a risk assessment to evaluate the FN curves for all the identified threats. The risk assessment could further take into account specific geographic areas, specific lines, specific installation dates, or even specific segments, in order to apply the ALARP framework to a more granular level of assets.

Due to the missing records, the operator would have to perform risk assessment using the most reasonable conservative assumptions to establish the pre-mitigation FN curves. From the FN curves, the operator would determine which region the FN curves fall in. If any FN curve falls in the intolerable region, the operator would have to devise the most cost-effective strategy to mitigate the risks until the FN curves fall into the ALARP region.

The most cost-effective strategy may not necessarily be wholesale replacement of the unknown polyethylene pipelines. Simply because an FN curve is in the intolerable region does not automatically mean the operator has to remove the pipelines. Although the operator has an absolute and immediate duty to drive the FN curve away from the intolerable region, the operator still has the leeway and the duty to evaluate its options for mitigation to select the most cost-effective option. It may be that immediate wholesale replacement of pipelines is not the best option.

Once the operator has driven the FN curve away from the intolerable region, the FN curve is now in the ALARP region. The operator will now have to evaluate its options again to devise the most cost-effective strategy to apply mitigation until the spending on mitigation produces a cost/benefit ratio that exceeds the gross disproportionality ratio or until the post-mitigation FN curves fall into the broadly acceptable region, whichever condition is met first.

4. ALARP CRITERION OF GROSS DISPROPORTIONALITY

Demonstration of gross disproportionality can be met by a cost/benefit ratio test.

Conceptually, gross disproportionality exists when the cost of risk mitigation greatly exceeds the benefit derived from the mitigation activity. These quantities may be defined either incrementally or cumulatively. The incremental approach defines the C/B ratio as the incremental costs divided by the incremental benefits of a group of risk control measures targeting a specific threat. The cumulative approach defines the C/B ratio as the cumulative costs divided by the cumulative benefits.

As an example, take the case of gas pipeline leak survey. In a cumulative approach, all costs associated with leak survey could be counted once the pipe has been put into service and accumulated up to the present. The benefits would be the avoided cost to replace leaky pipes, the avoided monetized values of human lives saved, injuries prevented, and properties saved between the time the pipe was put into service and the present. This could be further refined to take into account the time value of money to bring all costs to the present. In an incremental approach, the costs and benefits would be counted for any yearly cycle or any rate case cycle, or any period selected by the operator, for example.

If an FN curve starts from the intolerable region (Case 1, Figure 8), the cost to bring it down to the ALARP region should not be included in the cost/benefit ratio test since an operator has an absolute duty to reduce risk regardless of cost until the FN curve falls below the intolerable region. Both the cumulative approach and the incremental approach should only count costs and benefits **after** sufficient risk mitigation activity has pushed the FN curve to just below the upper tolerability limit line in the ALARP diagram (see Figure 6) since the concept of gross disproportionality only makes sense in the ALARP region.

There are advantages and disadvantages to each approach.

The incremental definition has an advantage in being intuitively clearer than the cumulative definition, since the incremental concept more closely measures the changing marginal effectiveness of risk mitigation spending. The primary disadvantage of the incremental definition is that the C/B ratio may vary wildly as mitigation expenditure is increased from period to period and the corresponding benefits may behave not in a proportional fashion. Economy of scale, cost inflation, and increasing efficiency/productivity due to additional experience are some possible factors among many others that may explain why incremental cost and incremental benefit may not necessarily increase or decrease in direct proportion to each another. As a result, the C/B ratio calculated using the incremental basis may not even follow a continuous curve from period to period as expenditure is increased.

With a cumulative definition, the C/B ratio will follow a smoothly continuous curve since the volatility of the ratio will be greatly moderated by the summation process. The primary disadvantage of the cumulative definition is that it could prematurely excuse an operator from further duty to mitigate risks if the operator had front-loaded mitigation activities (and expenditures) during the early phases of a utility's asset life or risk-causing activity. This undesirable effect would be particularly problematic for mitigation activities that did not result in enduring benefits that accrued to future periods.

For example, there would be a diminishing return of risk reduction against increasing repetition for many routine maintenance activities such as inspections and foot patrols of facilities. Repeating routine maintenance work on the same facility one day after the initial maintenance work has been completed would yield virtually no additional risk reduction but would incur twice the maintenance cost. Suppose during the early phase of the facility's life the operator had scheduled numerous such repeat maintenance works, the numerous repetitions would yield little additional cumulative risk reduction benefit that would accrue to future periods. However, the cost for each maintenance task would still count toward the cumulative total. With the incremental approach, the cost and benefit totals would reset at the start of a new

incremental period and would not experience this problem with front-loading. With the cumulative approach, any problem with front-loading would persist throughout the asset's life.

Additionally, the cumulative approach may require consideration of the time value of money (i.e. the concept that a dollar today is worth more than a dollar in the future) and the proper discount rate in calculating costs and benefits. Discounting of benefits based on the time value of money introduces another layer of controversy since the effect of discounting is to value a life saved today more than a life saved in the distant future.

For these reasons, the best approach may be a combination of the incremental definition and the cumulative definition. For example, expenditure and benefit figures for annual (biennial, triennial, or even longer periods) may be used to calculate the C/B ratios.

However, even in the UK where the ALARP concept originated there is no definite specification of what constitutes gross disproportionality or what that threshold disproportionality cost/benefit ratio should be. For example, guidelines created by the Office of Nuclear Regulation of the HSE give a suggested range for cost/benefit ratios when gross disproportionality is reached as between 2 and 10, depending on the severity of risk.²⁴ In other words, the costs of addressing a risk could range between twice the expected benefit and ten times before considering it disproportionate – quite a range, with little sense of where to draw the line. The HSE's R2P2 document likewise provides no guidance on what precisely constitutes gross disproportionality.

Likewise, the U.S. Army Corps of Engineers does not specify a precise gross-disproportionality ratio and states, "What is gross disproportion when determining ALARP is a matter for judgment." The Army Corps of Engineers goes on to list a suggested gross disproportionality

²⁴ HSE Office of Nuclear Regulation Guidance on the Demonstration of ALARP (As Low As Reasonably Practicable), Document ID: NS-TAST-GD-005 Revision 6.

http://www.onr.org.uk/operational/tech_asst_guides/ns-tast-gd-005.pdf

cost/benefit ratio in the ALARP region as between 2 and 10. That is, the cost of mitigation could be between two times and ten times the calculated benefit.

In essence, as shown by the HSE and the U.S. Army Corps of Engineers' examples, there can be much subjectivity and interpretation in determining what constitutes gross disproportionality since the test for gross disproportionality often boils down to selecting by the duty-holder a ratio that can vary widely from case to case.

In a workshop held at the CPUC on December 4, 2015 to discuss a preliminary version of this paper, some intervenor parties expressed concern that application of ALARP and the gross disproportionality concept would drive up utility rates. In actuality ALARP is meant to help balance safety with cost (utility rates). One could in fact argue that applying ALARP with optimization could conceivably lower utility rates since the operator would no longer be required to keep on spending money to mitigate a risk once the ALARP criterion of gross disproportionality has been met for that risk. For the many risks that are in the broadly acceptable region, the ALARP framework would allow the operator the latitude not to spend any money on mitigation, subject to compliance with other existing regulatory requirements. This could be another justifiable area for reducing existing spending and possibly lowering utility rates.

For those very rare risks that are in the intolerable region, the operator has an absolute duty to reduce the risks until they fall away from the intolerable region. This is after all the very definition of intolerable region. A potential nuclear disaster is a perfect example of an intolerable risk. It is exceedingly rare and is intolerable, but that does not mean the operator should not exercise prudence to select the most cost-effective strategies to mitigate an intolerable risk. The vast majority of risks are either in the ALARP region or broadly tolerable region. For risks in the ALARP region, the operator still has to select the most cost-effective mitigation strategy.

Whether an explicit gross disproportionality ratio should be set by a regulatory body and what that ratio (or range of ratios) should be are questions that can only be answered by a regulatory body in charge of the rate case proceedings. The lack of a definite gross disproportionality ratio does not detract from the basic theoretical soundness of the ALARP framework, but its absence could become a catalyst for contentiousness in revenue requirement decisions in rate cases. In such a scenario, one could possibly argue that having ALARP framework would be no better than the current status of evaluating rate cases without an ALARP framework, with all its uncertainties and contentiousness. What is certain, however, is that even in the absence of a definite and explicit gross disproportionality ratio an ALARP risk-informed decision framework would still offer the advantage of using a systematic and structured method to help force the operator and all stakeholders to recognize the tradeoff among safety, reliability, and utility rate affordability and to frame their conversations in these terms.

5. HOW DOES ALARP DIFFER FROM TRADITIONAL COST/BENEFIT ANALYSIS?

Traditional cost/benefit analysis (CBA) relies on a simple go/no-go criterion of whether the cost exceeds the benefit, i.e. if $C/B > 1$, then stop. In other words, in traditional CBA when applied to utility safety programs, the test criterion dictates that risk mitigation should (or could)²⁵ stop whenever costs exceeds benefits. Under traditional CBA, if the projected cost exceeds the projected benefit, then no mitigation is even necessary to begin with.

In ALARP, cost/benefit analysis is still used, but the analysis is applied only in the ALARP region. Furthermore, when it is applied in the ALARP region, the criterion is now: if cost disproportionally outweighs the benefit, then stop, i.e. if $C/B > \text{gross disproportionality ratio}$, then stop. In other words, mitigation would proceed and continue until cost disproportionately

²⁵ The choice between using “should stop” vs. “could stop” in traditional CBA rests on whether the emphasis of the CBA framework is to save resources (in the case of “should stop”) or to have the option to go beyond minimum safety (in the case of “could stop”).

exceeds benefit. A regulatory body is expected to promulgate a gross disproportionality ratio that is much greater than 1.

A natural question one might ask is: Why not dispense with the ALARP framework by simply using $C/B >$ gross disproportionality as the criterion in all cases?

The answer is that an ALARP framework is much more than simply using a traditional C/B ratio test of comparing cost against benefit. In the traditional approach one looks at the aggregate effect on risk from all threats, then identifies the mitigation measures that have the most cost beneficial impact on risk, and decides whether implementation of these measures will drive the risk below the upper tolerability limit.

Under the ALARP framework, one first looks at each threat to see where it stands in relation to the tolerability thresholds by examining its risk characteristics (as shown by its FN curve) to determine whether the C/B ratio test is even necessary. Blind application of the C/B ratio tests without the ALARP framework (and the examination of the FN curves in relation to the tolerability limits) would create unnecessary mitigation for threats that have very low risks. Lastly, the tolerability limits with the downward sloping lines in Figure 2 capture society's increasing aversion to risk as the severity of an accident increases. This is something a simple C/B analysis alone cannot capture.

The ALARP framework requires the operator to continually monitor changes in the character of the threat as embodied in the changing shape and location of its FN curve. This result is then fed back to the decision as to whether a C/B ratio test will be necessary.

The ALARP framework helps to filter out threats with low risks where no mitigation is needed. It then helps to address the question of how much mitigation is needed for those threats in the ALARP region. Combined with an optimization scheme, ALARP allows an operator to obtain the optimal combination of risk mitigation measures to reduce the most risk at the lowest cost. This, to reiterate, is something a simple $C/B > 1$ test by itself cannot achieve.

6. VALUE OF A STATISTICAL LIFE

For those unfamiliar with ALARP, the term often elicits an unpleasant association with “placing a monetary value on human life.” This negative reaction arises from the cost/benefit ratio test in ALARP, which requires the expression of costs and benefits to be in a common unit of measurement. Since costs are in monetary terms, this process in effect also requires the translation of predicted benefits (reductions in loss of life, reductions in injuries, and reductions in property damage, etc.) into monetary terms.

The unpleasant association of ALARP with monetizing human lives is unfounded, since traditional CBA often places a value on human life in exactly the same manner and yet traditional CBA elicits no such negative reaction. Traditional CBA in fact enjoys wide acceptance and is widely practiced in all walks of life, including usage at U.S. federal government agencies.^{26,27} The primary difference is that while traditional CBA is applied to many activities in all walks of life, including those activities that can affect human lives, ALARP is applied exclusively to those activities that can affect human lives. However, in those areas where traditional CBA is applied to activities that can affect human lives, translating the impact on human lives into monetary terms is a readily accepted aspect of traditional CBA.

The conventional way to measure the benefit of reductions in loss of life relies on a concept referred to as the Value of a Statistical Life. A similar term used in risk management referring to

²⁶ Executive Order 12866 (58 FR 51735; October 4, 1993), E.O. 12866

<http://www.archives.gov/federal-register/executive-orders/pdf/12866.pdf>

Under, E.O. 12866, federal agencies must assess all costs and benefits of available regulatory alternatives.

²⁷ Executive Order 13563 (76 FR 3821; January 21, 2011), E.O. 12866

<http://www.gpo.gov/fdsys/pkg/FR-2011-01-21/pdf/2011-1385.pdf>

E.O. 13563 reaffirms and amplifies the principles embodied in E.O. 12866. E.O. 13563 requires federal agencies to quantify anticipated benefits and costs of proposed rulemakings as accurately as possible using the best available techniques, and to ensure that any scientific and technological information or processes used to support their regulatory actions are objective.

the same concept is the Implied Cost of Averting a Fatality (ICAF). In the words of the U.S. Department of Transportation's (USDOT) 2014 Guidance on Treatment of the Economic Value of a Statistical Life (VSL), it is defined as a measure of the "additional cost that individuals [or society] would be willing to bear for improvements in safety (that is, reductions in risks) that, in the aggregate, reduce the number of fatalities by one."²⁸

As pointed out in the USDOT guidance, "what is involved is not the valuation of life as such, but the valuation of reductions in risks." This nuanced distinction (between valuation of life vs. valuation of reductions in risks) may well be nothing more than academic, but what is important is that VSL measures the benefit of risk reduction from a societal perspective rather than from an individual perspective. Different individuals have different perceptions of what risk reduction is worth particularly as the benefit is accrued to the individual, but from a societal point of view, variability in individual perceptions does not come into play. There are however different methods of measuring the economic value of statistical life with the various methods giving significantly different estimates of the VSL.

Whereas a precise estimate of VSL is of great importance in traditional CBA, under ALARP having a consistent VSL for all operators and all risks (whatever that adopted value may be) is far more important than the value itself. The reason for this disparity lies in the fact that under ALARP gross disproportionality is specified by both the selected VSL and the selected gross disproportionality ratio. In adopting an ALARP framework, the regulatory body would select both a VSL and a gross disproportionality ratio in tandem. If the regulatory body adopted a smaller than justified VSL, the effect would tend to be counteracted by the selection of a larger than expected gross disproportionality ratio, and vice versa. The reason for this is that as long as a risk is in the ALARP region, the question for a regulatory body is no longer purely about safety but rather about deciding the tradeoff between safety and rate affordability. Referring

A detailed analysis and history of Executive Orders, legislations, and court cases addressing cost-benefit analysis is found in Alison's Prout's PUTTING A PRICETAG ON LIFE: THE VALUE OF LIFE AND THE FDA, <https://dash.harvard.edu/handle/1/8848247>.

²⁸ Guidance on Treatment of the Economic Value of a Statistical Life (VSL) in U.S. Department of Transportation Analyses - 2014 Adjustment. http://www.transportation.gov/sites/dot.gov/files/docs/VSL_Guidance_2014.pdf

to Figure 7, we can see that the combined effect of the selected VSL and gross disproportionality ratio is to determine how far down we want to push the FN curve (i.e higher safety) before mitigation should stop.

Since there is no magic number that signifies the “correct” disproportionality ratio, the regulatory body would tend to select a VSL and a gross disproportionality ratio in concert to result in some affordable level of risk mitigation expenditure in the ALARP region. In attempting to strike some balance between safety and rate affordability, the regulatory body would tend to choose the VSL and gross disproportionality ratio in such a way that the “error” in one would tend to be moderated by the “error” in the other. Traditional CBA, with the C/B ratio fixed at 1, enjoys no such leeway. With traditional CBA, the selected VSL has a direct and unmitigated impact on decision-making. These observations should not be misconstrued to imply the insignificance of the proper selection of an applicable VSL under ALARP. The proper selection of VSL is still important under ALARP, just not nearly as critical as it is under traditional CBA.

This paper is not about the methodologies of estimating VSL, and since there are myriad publications on the subject of VSL, a deep dive on VSL estimates is beyond the scope of this paper. We instead list in the footnotes and bibliography a couple of very good publications analyzing the value of life.^{29,30} If and when that bridge needs to be crossed in the future when an ALARP framework is to be seriously contemplated by a regulatory body, this topic can be addressed in much greater depth at that point.

The Department of Transportation’s (USDOT) 2014 Guidance on Treatment of the Economic Value of a Statistical Life (VSL) is a possible candidate for consideration. The 2014 revision of the DOT Guidance identifies \$9.2 million as the best point estimate of VSL for base year 2013, with a low end estimate of \$5.2 million and a high end estimate of \$13.0 million.

²⁹ Alison Prout, PUTTING A PRICETAG ON LIFE: THE VALUE OF LIFE AND THE FDA
<https://dash.harvard.edu/handle/1/8848247>

³⁰ W. Kip Viscusi, Fatal Tradeoffs (1992)

The DOT Guidance also presents a formula to estimate the VSL for future years based on the equation:

$$VSL_{2013+N} = VSL_{2013} \times 1.0118^N$$

where VSL_{2013+N} is the VSL value N years after 2013.

The DOT Guidance also lists a table to show the relative equivalent fraction of VSL for various levels of injury:

AIS Level³¹	Severity	Fraction of VSL
AIS 1	Minor	0.003
AIS 2	Moderate	0.047
AIS 3	Serious	0.105
AIS 4	Severe	0.266
AIS 5	Critical	0.593
AIS 6	Unsurvivable	1.000

This paper has so far considered the benefits of risk mitigation exclusively in terms of the monetized avoided costs of deaths and injuries. A variation of the basic ALARP concept could be easily broadened to include avoided costs of property damage, monetized avoided costs of environmental damage, monetized avoided costs of outage (this would be a way to broaden the ALARP concept to take into account reliability), plus other applicable avoided costs. These additional dimensions would be highly relevant in risk mitigations dealing with wildfires.

Yet another variation of the basic ALARP concept is to express the FN curves and the ALARP diagram in terms of dollars instead of number of fatalities. In this variation the horizontal axis would be dollars and the vertical axis would be the annual probability of exceeding the

³¹ AIS stands for Abbreviated Injury Scale.

monetized avoided costs. This variation would complement the variation to include avoided costs of property damage, environment damage, and outage.

7. HOW DOES AN ALARP FRAMEWORK RELATE TO RISK TOLERANCE IN ENTERPRISE RISK MANAGEMENT?

How does risk tolerance fit in the total picture of ALARP? Risk tolerance can be defined as the maximum amount of residual risk that an entity and its stakeholders are willing to accept after application of risk control measures.³² There are three relevant perspectives to risk tolerance: societal perspective, regulatory perspective, and utility operator (corporate) perspective. Risk tolerance from utility employees' perspective is subsumed in one or more of the other perspectives.

In regulatory applications, the ALARP framework typically uses societal risk tolerance as the relevant risk tolerance in risk mitigation activities. Societal risk tolerance is represented by the risk tolerance limit lines in the ALARP diagram. Societal risk tolerance and regulatory risk tolerance should coincide except in two crucial respects: a regulatory body has the duty not only to protect the public from physical harm, but also 1) to ensure the economic viability of the regulated utilities under its jurisdiction and 2) to balance safety (and reliability) with rate affordability. For the purpose of this paper, we can assume there is no distinction between societal risk tolerance and regulatory risk tolerance because: 1) the effect of the first difference would only come into play in extraordinarily rare instances where the economic viability of the utility might be threatened if it were required to comply with the full ALARP framework for all identified risks, and 2) consideration of rate affordability is baked-in and implied in the ALARP framework through application of the ALARP cost/benefit test.

³² This definition is similar to that used by FERMA, Guidance on the 8th EU Company Law Directive, article 41, Part 2, 2011 (<http://www.ferma.eu/app/uploads/2011/12/ecia-ferma-guidance-on-the-8th-eu-company-law-directive-part-2.pdf>), which defines risk tolerance as: "The maximum amount of risk that the company can bear despite controls." Other definitions of risk tolerance can be found in RIMS Executive Report, The Risk Perspective, Exploring Risk Appetite and Risk Tolerance (https://www.rims.org/resources/ERM/Documents/RIMS_Exploring_Risk_Appetite_Risk_Tolerance_0412.pdf).

For the rare events where a utility's viability might be threatened, the operator could seek relief from the regulatory body to use a reduced gross disproportionality ratio to address those rare specific risks that threaten the economic viability of the operator. The regulatory body's ability to grant such deviations would in effect preserve the equality of societal risk tolerance and regulatory risk tolerance.

Next, we turn our attention to risk tolerance from an operator's perspective. Since a utility operator's foremost instinct is to mitigate risks that can threaten the viability of the company, the risk tolerance from a utility operator's perspective is that which is embodied in its Enterprise Risk Management (ERM) program. Under ERM, risk tolerance can be defined as the maximum amount of residual risk that an organization is willing to accept after application of risk control measures. A clearly stated risk tolerance level should be part of any well-designed risk management system. An informal relationship can be shown to exist between safety-related component of risk tolerance (under ERM) and societal risk tolerance (i.e. the upper or lower boundary risk tolerance lines under ALARP) through the arguments below.

We first recognize that ERM might consider risk from a multitude of components, for example, financial risk, risk from safety-related events, risk to the environment from the company's actions, risk to enterprise goodwill and reputation, to name just a few possible components. Risk tolerance may furthermore be set for the whole enterprise or separately for each line of business within the organization.

For a public utility, one would expect risk arising from safety-related events involving the utility's facilities to be the primary contributor to total enterprise risk. Safety-related events should, therefore, be the primary driver of enterprise risk tolerance for a public utility. The ALARP framework, on the other hand, is only concerned with safety-related risks. In order to illustrate the relationship between safety-related ERM risk tolerance and ALARP risk tolerance, we make the simplifying assumption that the ERM safety-related risk tolerance and the ALARP tolerance limits share identical physical dimensions (or units) of specification.

In the long run, no prudent utility operator could mitigate a system's safety-related risks down to the upper threshold line in the ALARP diagram and decide to go no further without facing serious regulatory (and also probably legal) repercussions. This is because, by definition, the upper tolerability threshold is chosen to represent the most safety-related risk that society (and a regulatory body) would tolerate. Society (and a regulatory body) would expect an operator to keep on mitigating risks as long as the costs are affordable and the benefits could justify the costs. Likewise, in the long run, an operator would face stiff opposition if it attempted to mitigate risks to such a low level that the mitigation measures would translate to prohibitively high utility rates. For these two reasons, an operator would either intuitively, or be forced through regulatory rate-setting proceedings, to operate in a de facto regime similar to the ALARP concept. The upper and lower tolerance values in ALARP can then be viewed as the band of uncertainty surrounding the ERM safety-related risk tolerance due to underlying regulatory rate-setting actions.

In summary, if adopted an ALARP framework would explicitly express what is intuitively practiced by utility operators and regulators. An operator would intuitively mitigate risks until the expected frequency of occurrence falls between the two tolerability limits in the ALARP diagram.

The next section describes the relationship between an ALARP framework at the enterprise level and line of business level and that at the threat/risk level.

8. ALARP FRAMEWORK AT THE THREAT AND RISK LEVEL

The potential usefulness of an ALARP framework in public utility rate cases lies in its application to guide risk mitigation decisions at the programs and projects level. One should recognize at the outset that an ALARP framework to utility safety does not by itself provide answers to what mixture of mitigation measures a utility should employ or what pace of execution to apply

those selected mitigation measures. An ALARP framework can only provide an answer to the question of whether additional risk mitigation is needed for each identified threat.

The forgoing statements are true whether the ALARP framework is applied to the enterprise level, line of business level, or threat/risk level. A complication may arise in which risk mitigation is continually exercised as part of a maintenance program. An ALARP framework cannot generally by itself distinguish between this routine maintenance activity and any additional mitigation activity needed to satisfy the ALARP criterion so long as the same amount of mitigation work is done in the same period. Optimization techniques would be needed to help select the pace, intensity, or mix of risk mitigation strategies.

As pointed out earlier, the upper and lower tolerability limit lines in the ALARP diagram are values set by a regulatory body and not based mathematical derivation. The same observation can be made that the tolerability limits at the threat level are also promulgated values. A good starting point for the selected threat level tolerability limits is that they should be equal to the line of business level tolerability limits divided by the number of identified threats. This starting point is based on the observation that accident frequencies at the enterprise levels are composed of the sum of component threat level frequencies. This is generally a good approximation when threats are considered independent. There are interacting effects among the threats, but their contribution to the overall enterprise level frequency is expected to be small. For this reason, the promulgated values for threat level tolerability limits can be expressed as:

Upper tolerability limit at threat level = Upper tolerability limit at line of business/ n ,

Lower tolerability limit at threat level = Lower tolerability limit at line of business/ n ,

Where n = number of all identified threats.

A separate FN curve is derived for each identified threat based on either empirical data or loss distribution model based on empirical data. Again, there is interaction among certain threats, but we will assume independence of threats for the purpose of applying FN curves to the ALARP

diagrams at the threat level. The same criteria described in earlier sections about FN curves superimposed on an ALARP diagram would then apply for each threat. This then would form the basis of risk mitigation needed and when mitigation can cease at the program and project level.

9. AN ALARP FRAMEWORK AMONG COMPETING OBJECTIVES AND CONSTRAINTS

Without an ALARP framework, decisions on what mixture of mitigation measures to use or how quickly to apply those selected mitigation measures are made through subjective judgment informed by best practices, compliance requirements, availability of resources, etc. With the ALARP specifications at the threat level, it becomes possible to use optimization routines to arrive at some optimum mixture of programs and projects, and application pace of these programs and projects, subject to various constraints, to maximize safety and reliability and concurrently to minimize rate increase. The ALARP specifications and the ERM risk tolerance levels at both the enterprise level and the line of business level would be simultaneous constraints in this optimization scheme.

An ALARP framework in conjunction with optimization routines could lead to an optimal selection of mitigation programs and projects to achieve either:

- 1) The lowest levels of combined residual risks (i.e. maximum combined risk reduction) at a targeted level of revenue requirement and rate increase, or
- 2) The lowest level of revenue requirement and rate increase for a targeted level of combined risk reduction, or
- 3) An optimal balance of low residual risks and low revenue requirements.³³

³³ There could be multiple optimal solutions in an optimization routine, with all such optimal solutions being equally desirable.

Without such optimization, an operator could only be guided by experience, subjective judgment, safety regulations, best practices, resource constraints, etc., with no assurance that whatever the operator puts forth in a rate case request would be an optimal proposal.

Such optimization schemes may fail in a regulatory environment where an unrealistic goal of “no accidents ever allowed” is imposed. If such a condition was imposed, either an infinite amount of safety spending would be required or a solution that would contradict the original “no accidents ever allowed” specification would be produced. Mathematically speaking, every risk, no matter how improbable, has a non-zero potential to cause an accident (unless the activity associated with that risk has been disallowed altogether).

Mitigation schemes can only reduce the risk arising from a threat to a very small, but still finite, non-zero number. For example, for some particular threat, mitigation may drive the pre-mitigation expected frequency from one accident per year to a post-mitigation expected frequency of one accident every thousand years, but it will still be some finite number, however small that expected frequency number may be. The problem with the “no accidents ever allowed” expectation in conjunction with the expectation of rate affordability fundamentally boils down to over-specifying a problem with conflicting and irreconcilable objectives.

An ALARP framework would be a useful tool to help demonstrate to the regulatory body and parties involved in a rate case, in an objective and quantifiable way, the unaffordable rate increase implications that would result from such an unrealistic expectation.

10. BEST PRACTICES IN AN ALARP FRAMEWORK

The term “best practices” is a business buzzword generally referring to a set of operational practices that organizations adopt to produce superior results. As applied to utility safety, the term refers to a set of operation and maintenance procedures, operations, or standards that

utility operators as a group have generally accepted to be conducive to operating and maintaining their systems in a safe manner. Best practices are at least as stringent as minimum regulatory code requirements at the national level. It is possible, however, for best practices to be less stringent than state safety codes, since individual states may have adopted more stringent requirements than the national codes.

Where legislation requires the application of “best practices”,³⁴ an operator using an ALARP framework to utility safety would select the more stringent of the requirements dictated by either ALARP or best practices. For example, suppose industry best practice dictates that leak surveys be performed at a certain frequency, for example, every 12 months. In the absence of an ALARP framework, this 12-month frequency would form the baseline of minimum leak survey activity. However, suppose that under an ALARP framework with optimization, an operator determines that leak surveys should be performed every 15 months. Under this scenario in an ALARP framework, the operator should perform leak survey based on the more stringent of the requirements, i.e. in this case the operator should perform leak survey every 12 months.

Suppose on the other hand that ALARP with optimization determines that leak surveys should be performed every 10 months. In this case since a 10-month frequency is more stringent than the 12 months dictated by best practices, the operator should perform leak survey every 10-months, in the absence of other considerations. (In actuality, industry best practices that the operator chose to adopt could be incorporated as additional constraints in the optimization routines).

When industry best practices are **not** required by legislation, the choice between the levels of activity dictated by industry best practices and those indicated by ALARP with optimization may be less clear. On the one hand, an operator may be reluctant to go below industry best practices for fear of being found negligent were lawsuits to result from an accident. On the

³⁴ For example, California’s 2011 Senate Bill 705:
http://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201120120SB705.

other hand, industry best practice may reduce the level of risk for that threat to a level beyond necessary as determined by ALARP. If the latter condition were to happen, there would be a de facto waste of resources as the additional spending on risk mitigation above that dictated by ALARP could have been better spent elsewhere to more effectively reduce total system risk from other threats.

This illustration points to the need to use ALARP with optimization to look at overall system risk from all threats and not just an individual risk and individual best practice in isolation. Furthermore, if an ALARP framework was formally adopted into regulation, the regulated utility would have sufficient legal justification to deviate from industry best practices when an ALARP framework indicates it to be reasonable to do so. This should allay any concern a utility may have of lawsuits accusing it of negligence for failing to follow industry best practices.

11. HURDLES TO ACCEPTANCE OF AN ALARP FRAMEWORK

One may reasonably envision several prominent hurdles to wide acceptance of an ALARP approach to public utility risk-based decision making in the United States:

1. Need to place value on statistical life - This hurdle has already been addressed in Section 6 (Value of a Statistical Life.)
2. Unfamiliarity - Before regulators and stakeholders would embrace ALARP, they would first have to understand the details of how ALARP works. This may prove to be a challenge for those not well versed in the technical concepts of probability and statistics.
3. Lack of data and deficiencies in quantitative models – Although the cost side of cost/benefit analysis may be straightforward to determine, estimation of benefit would require quantitative probabilistic risk assessment models to estimate expected accidents avoided during any period of time due to the incremental risk mitigation controls. Data needed to model such loss

scenarios may be lacking and/or unreliable and the models may be lacking in rigor with inadequate predictive power. This is not an impassible hurdle since missing data can be estimated by subject matter experts (SMEs).

4. ALARP requires an explicit statement of risk tolerance, stated in the form of the upper and lower risk tolerability limits in the ALARP diagrams. This may well be the biggest potential hurdle of all. A cultural change would need to take place to inculcate the idea that risk can never be driven down to zero with any affordable safety budget in the real world and that society must be willing to accept an explicit residual risk as a tradeoff between safety and utility rate affordability.

5. Absent more robust quantitative risk models, it would be difficult to attribute an incremental benefit to any particular threat and any particular mitigation control measure associated with that threat. For example, leak survey is one type of risk mitigation but the total benefit must somehow be apportioned to all threats that have benefited from the leak survey, including for example, internal corrosion, external corrosion, stress corrosion cracking, earth movement, etc. Is there a technically sound way to apportion the total incremental benefit to all threats in the models? It is an unanswered question.

6. Regulators' "no accidents" expectation does not conform with reality - One possible way to reconcile reality with expectation is to analyze what is actually meant by such absolute statements as "our goal is zero accidents" or "no accidents allowed," or "zero tolerance for accidents."

For example, the CPUC has a vision for "zero accidents" caused by the public utilities it regulates: "Ultimately we are striving to achieve a goal of zero accidents and injuries across all the utilities and businesses we regulate, and within our own workplace." ³⁵ What does it really

³⁵ Safety Policy Statement of the California Public Utilities Commission, http://www.cpuc.ca.gov/NR/rdonlyres/967047D4-19CE-45B1-8766-057F1D7FF1CD/0/VisionZero4Final621014_5_2.pdf

mean in reality? First, it should be pointed out that it is a vision and not a mandate or outright prohibition. A vision is a tacit acknowledgement that reality will usually turn out to be quite different from the vision for any number of reasons. Affordability of utility rates is one reason why vision and reality may not always coincide. As pointed out earlier, absolute zero risk means either an infinite safety budget or cessation of the risk-causing activity altogether. The laws of probability are another reason why reality and vision may diverge. A law could be passed that proclaims the sun is forbidden from rising tomorrow, but the proclamation would violate the natural laws of the heavens, and nature would simply defy it. The laws of probability are incompatible with a “no accidents and injuries (ever allowed)” and nature will simply violate such an absolute proclamation.

However, let’s take a step back to try to understand what the CPUC regulators actually meant when they proclaimed “Ultimately we are striving to achieve a goal of zero accidents and injuries across all the utilities and businesses we regulate, and within our own workplace.” It means the CPUC expects a regulated utility to continually improve safety and provide service in the safest way that it possibly can within the limited budget that it is authorized to collect from ratepayers. It means that despite not always being able to prevent accidents and injuries due to the invisible laws of nature and probability, the utilities should still strive for the best safety record possible. It means resources and efforts must be disproportionally expended to prevent accidents of the types that cause serious injuries and fatalities. It does not mean that the utility should spend its entire safety budget to prevent superficial injuries such as shallow skin punctures from wooden power pole splinters by replacing all wooden poles that are still reliable. It expects a utility to allocate its resources to balance cost against the frequency and severity of potential risks. That is what “vision zero” really means.

When reasoned out this way, it becomes plausible that perhaps the CPUC’s “vision zero” policy is not so incompatible with reality after all. In fact, when we reason it out in this nuanced fashion, a safety and risk management framework based on ALARP suddenly looks entirely compatible with the “vision zero” expectation.

12. CONCLUSION: WHY ADOPT AN ALARP FRAMEWORK IN RATE CASES?

This paper has highlighted several reasons why an ALARP framework applied to utility risk management and rate cases could offer some important advantages:

1. Introduction of an ALARP decision framework would provide a strong driving force for bolstering a utility's incentive for strengthening its risk assessment capability and risk-informed decision process.
2. Coupled with an optimization scheme subject to constraints, an ALARP framework could allow an operator to determine an optimum mix of risk mitigation measures, executed at the optimal pace, and to the optimal extent to maximize safety and reliability in the most cost effect manner.
3. An ALARP framework would offer operators a consistent and definite regulatory signal on the expected balance between utility safety and rate affordability.
4. An ALARP framework would be a useful tool to help demonstrate to the regulatory body and stakeholders in a rate case, in an objective and quantifiable way, the tradeoff between safety and utility rates.
5. Adoption of an ALARP framework would overcome utility operators' reluctance to state an explicit risk tolerance in their enterprise risk management. A regulatory body's promulgated upper and lower tolerability limits in ALARP would act as the equivalent risk tolerance limits in enterprise risk management.

13. RECOMMENDATIONS AND ROADMAP AND REQUIRED ELEMENTS OF ALARP IMPLEMENTATION

This paper explored the mechanics and benefits of the ALARP risk management principle with the goal of providing sufficient details to help ignite a discussion among stakeholders. If applied

correctly, the potential benefits of an ALARP framework in utility rate cases are huge, but in many ways the hurdles to adoption of ALARP are also huge. None of the hurdles listed in Section 11 (Hurdles to Acceptance of an ALARP Framework) are unsurmountable. The solution to each one of the hurdles is simply more time, time needed to familiarize the regulators and stakeholders of the concept, time needed to bring about a wide acceptance of an explicit risk tolerance, of the fact that risk can never be driven to zero, and of the absolute tradeoff between safety and utility rate affordability. However, the discussion to bring about the application of ALARP to improve safety and utility rate cases should not wait until the time has passed to topple the hurdles.

Regulators and stakeholders in utility rate cases should start the discussion now to consider implementing ALARP into rate cases. The implementation of an ALARP risk management framework would at a minimum require the following basic steps:

1. Regulatory entity specifies the elements of an ALARP diagram by promulgating the upper and lower tolerability limits (at $N=1$) for enterprise level, line of business level, or threat level and the slope of the tolerability limit lines. This is in effect adoption of explicit levels of risk tolerance.
2. Regulatory entity adopts the value of a statistical life by reference to well-known published values. This value will update based on the updates of the referenced value.
3. Regulatory entity either promulgates a gross disproportionality ratio (or a range of ratios) to be used in the ALARP test or provides language in the regulatory framework to allow an operator to select an applicable disproportionality ratio from a range based on its own circumstances that it must justify in front of the regulatory entity.
4. Utility operators continue to build loss models and collect data with an eye toward eventual full use of quantitative risk assessment and ALARP model. Utility operators should consider direct sharing loss data to enhance data credibility.

5. Utility operators will initially file rate cases based on relative risk scoring model. Utility operator will incrementally apply an ALARP based model as a check on the risk scoring model. As more data become available to permit fuller use of ALARP model, the ALARP framework will become more dominant in future rate cases.

14. BIBLIOGRAPHY

1. CPUC Decision 14-12-025, Decision Incorporating a Risk-Based Decision-Making Framework into the Rate Case Plan and Modifying Appendix A of Decision 07-07-004.
2. Excerpt from Edwards v. National Coal Board,
http://www.safetyphoto.co.uk/subsite/case%20e%20f%20g%20h/edwards_v_national_coal_board.htm
3. Excerpt from Edwards v. National Coal Board, Court of Appeal decision,
<http://www.hse.gov.uk/risk/theory/alarpglance.htm>
4. UK by the Health and Safety at Work etc. Act 1974 (HSWA)
<http://www.legislation.gov.uk/ukpga/1974/37/contents>
5. R M Pitblado, M Bardy, P Nalpanis, P Crossthwaite, K Molazemi, M Bekaert, V Raghunathan, Det Norske Veritas
International Comparison on the Application of Societal Risk Criteria
6. Australian National Offshore Petroleum Safety and Environmental Management Authority (NOPSEMA) Guidance Note on ALARP: <http://www.nopsema.gov.au/assets/Guidance-notes/N-04300-GN0166-ALARP.pdf>
7. Danish Energy Agency's Guidelines on Safety and Health Related Conditions on Offshore Installations, etc.
8. Hong Kong Advisory Council on the Environment,
http://www.epd.gov.hk/epd/english/boards/advisory_council/eia_paper026.html
9. Aven, Terje; Vinnem, Jan Erik and Vollen, Frank; "Perspectives on Risk Acceptance Criteria and Management for Offshore Applications - Application to a Development Project";
International Journal of Materials & Structural Reliability Vol.4, No.1, March 2006, pp 15-25.
10. Interim Tolerable Risk Guidelines for US Army Corps of Engineers Dams
http://www.nfrmp.us/TRG2010/docs/USSD_2009_Corps_Tolerable_RisK_Dam_Safety_10_Feb_09.pdf
11. U.S. Army Corps of Engineers, Best Practices in Dam and Levee Safety Risk Analysis - Public Risk Tolerance and Risk Guidelines.
12. 10CFR835 "Occupational Radiation Protection" regulating ALARA in DOE facilities; and 10CFR20 regulating ALARA facilities licensed by the NRC as prescribed by 10CFR part 50, and

Part 52 and Part 70 –ANS Local Section Address Eric P. Loewen 2011 that's the commercial nuclear power plants, food irradiation facilities, medical facilities that handle radiation sources.

13. HSE Guidance on ALARP Decisions in COMAH,
http://www.hse.gov.uk/foi/internalops/hid_circs/permissioning/spc_perm_37/
 14. A full discussion on FN curves can be found in this HSE publication: **Transport fatal accidents and FN-curves: 1967-2001**, <http://www.hse.gov.uk/research/rrpdf/rr073.pdf>
 15. HSE document, *Reducing risks, protecting people* (often referred to by the acronym, R2P2),
<http://www.hse.gov.uk/risk/theory/r2p2.pdf>
 16. HSE Office of Nuclear Regulation Guidance on the Demonstration of ALARP (As Low As Reasonably Practicable), Document ID: NS-TAST-GD-005 Revision 6.
http://www.onr.org.uk/operational/tech_asst_guides/ns-tast-gd-005.pdf
 17. Executive Order 12866 (58 FR 51735; October 4, 1993), E.O. 12866
<http://www.archives.gov/federal-register/executive-orders/pdf/12866.pdf>
Under, E.O. 12866, federal agencies must assess all costs and benefits of available regulatory alternatives.
 18. Executive Order 13563 (76 FR 3821; January 21, 2011), E.O. 12866
<http://www.gpo.gov/fdsys/pkg/FR-2011-01-21/pdf/2011-1385.pdf>
E.O. 13563 reaffirms and amplifies the principles embodied in E.O. 12866. E.O. 13563 requires federal agencies to quantify anticipated benefits and costs of proposed rulemakings as accurately as possible using the best available techniques, and to ensure that any scientific and technological information or processes used to support their regulatory actions are objective.
- A detailed analysis and history of Executive Orders, legislations, and court cases addressing cost-benefit analysis is found in Alison's Prout's PUTTING A PRICETAG ON LIFE: THE VALUE OF LIFE AND THE FDA, <https://dash.harvard.edu/handle/1/8848247>.
19. Guidance on Treatment of the Economic Value of a Statistical Life (VSL) in U.S. Department of Transportation Analyses - 2014 Adjustment.
http://www.transportation.gov/sites/dot.gov/files/docs/VSL_Guidance_2014.pdf
 20. Alison Prout, PUTTING A PRICETAG ON LIFE: THE VALUE OF LIFE AND THE FDA
<https://dash.harvard.edu/handle/1/8848247>
 21. W. Kip Viscusi, Fatal Tradeoffs (1992)
 22. FERMA, Guidance on the 8th EU Company Law Directive, article 41, Part 2, 2011
(<http://www.ferma.eu/app/uploads/2011/12/eciia-ferma-guidance-on-the-8th-eu-company-law-directive-part-2.pdf>),

23. RIMS Executive Report, The Risk Perspective, Exploring Risk Appetite and Risk Tolerance (https://www.rims.org/resources/ERM/Documents/RIMS_Exploring_Risk_Appetite_Risk_Tolerance_0412.pdf)

24. Safety Policy Statement of the California Public Utilities Commission, http://www.cpuc.ca.gov/NR/rdonlyres/967047D4-19CE-45B1-8766-057F1D7FF1CD/0/VisionZero4Final621014_5_2.pdf

25. California 2011 Senate Bill 705
http://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201120120SB705

(END OF ATTACHMENT 1)

ATTACHMENT 2

**BEFORE THE PUBLIC UTILITIES COMMISSION
OF THE STATE OF CALIFORNIA**

**COMBINED UTILITIES
S-MAP UNIFORMITY REPORT**

**Prepared For Discussion at December 4, 2015
S-MAP Workshop**

**SAN DIEGO GAS & ELECTRIC COMPANY
SOUTHERN CALIFORNIA GAS COMPANY
SOUTHERN CALIFORNIA EDISON COMPANY
PACIFIC GAS AND ELECTRIC COMPANY**

December 1, 2015

TABLE OF CONTENTS

1. Executive Summary	1
1.1. Objective	1
1.2. Approach	1
1.3. Accomplishments and Findings	1
2. Introduction	3
2.1. Background	3
2.2. Objective	3
2.3. Approach	4
2.4. Report Structure	5
3. The Utilities' Risk Management Frameworks	6
4. Areas of Commonality	7
5. Areas of Uniqueness	13
6. Areas for Future Consideration	14
7. Commitment	15

1. Executive Summary

1.1. Objective

The California Public Utilities Commission's (CPUC or Commission) Order Instituting a Rulemaking (OIR) decision and subsequent Safety Model Assessment Proceeding (S-MAP) led PG&E, SCE and SoCalGas/SDG&E (the "combined utilities") to review areas where the companies could adopt a common approach to elements of their risk frameworks and processes.

1.2. Approach

Two utility workshops were held to identify areas of common approach and areas where differences between companies made adopting a common approach either difficult or costly. Based on the results of the utility workshops, this Uniformity Report has been prepared for review and discussion at the December 4, 2015 S-MAP workshop.

1.3. Accomplishments and Findings¹

The utilities began each of the utility workshop discussions by recognizing that while there were differences in how the companies address risks, they share the same objective of addressing safety and security risks. The workshops allowed the utilities to work together on achieving common standards and gaining an understanding of why - in some areas - a unique approach was important to the individual utility.

Working together, the utilities agreed on uniform approaches including but not limited to:

- The definition of the various safety impact levels in their 7x7 matrices.
- The definition of the various frequency levels in their 7x7 matrices.
- The application of principles and guidelines of International Standard Organizations (ISO) 31000 and 55000 as well as the Cyclac 10-step model.
- Common risk categories (e.g., Wildfire, Gas Explosion, etc.)
- A common framework for evaluating risk reduction benefits of mitigations.

Areas where the companies believe it is appropriate to continue to use different approaches were also agreed to including:

- The use of different risk specific approaches (e.g., wildfire risk).
- The use of different risk scoring algorithms.

Additional benefits of the uniformity workshops included:

¹ A common risk lexicon had been agreed to by all of the S-MAP parties prior to the utility uniformity workshops.

- A commitment by each utility to continue open and transparent discussions regarding additional potential areas of uniformity.
- A commitment by each utility to share information on how top safety and security risks are being assessed, managed and mitigated.

2. Introduction

2.1. Background

In its December 2014 OIR decision regarding the development of a risk-based decision making framework, the Commission indicated that, among other things, the Commission wanted the S-MAP to address the development of common standards and approaches. Specifically, the CPUC found:

1. There were efficiency benefits if the utilities were able to use common standards.
2. That while it is premature to decide whether one risk assessment approach, model or methodology should be adopted, the S-MAP proceeding should determine whether uniform and common standards could be developed to assess, manage and mitigate risks inherent in each utility's operations.
3. There are differences between utilities' demographics and between their systems and, therefore, there needed to be some flexibility in how each utility addresses, manages and mitigates its risks.²

At Workshop #3 of the 2015 S-MAP proceeding, parties agreed that it would be beneficial if the utilities addressed which elements of their risk management frameworks and processes could reach congruence or become more similar.³

The utility uniformity workshops highlighted that, by working together, the utilities could identify areas where common approaches to various elements of their risk methods could be standardized. The utilities also recognized the frameworks each company has adopted to identify, assess and manage their risks are essentially the same. The work as a group allowed each company to understand that while there are differences in implementation, the same risk categories are being addressed; and, while there are implementation differences, the same objectives of highlighting and addressing safety risks are achieved. Additionally, the utilities agreed that while there may be differences in implementation, there is a commitment to learning from each other and being transparent as to the objectives and the results of the common standard discussions. Finally, there was agreement on continuing to work towards common approaches where appropriate.

2.2. Objective

The objectives of the utilities' efforts included the following:

² Decision 14-12-025, p.51.

³ Safety and Enforcement Division Risk Assessment Section Final Staff Workshop #3 Report on Staff Evaluation of Utility Risk Assessment Models and Risk Assessment Methodology Phase (RAMP) A.15-05-002, et al.

- Expand learning and understanding of the different risk management approaches and models used by each utility.
- Identify areas where there were opportunities to apply a uniform approach to elements of each company's risk methodology.
- Identify areas where the utilities have to maintain their uniqueness because of demographic or systems differences.⁴

2.3. Approach

On October 29 and 30, a two-day utility workshop was held at SDG&E's Century Park offices. The focus of the initial workshop was on understanding the frameworks used by each utility to assess, manage and mitigate risks, and to identify areas where common approaches were attainable. At the initial utility workshop, a preliminary discussion of the risk scoring algorithm used by each utility was held. The companies felt it would be productive if further work and analysis could be done to assess whether a common risk scoring algorithm could be adopted. A second workshop was held at PG&E's San Francisco offices on November 6 to review and discuss the risk scoring algorithms. After the November 6 utility workshop, each company analyzed how the adoption of a common algorithm might affect their individual risk prioritizations.

In advance of the initial utility workshop, each company addressed three questions regarding their risk assessment, management and mitigation processes. The three questions can be summarized as follows:

1. Based on what has been discussed and your company's views, what areas do you believe are obvious areas for a uniform approach?
2. From your company's perspective, are there additional areas where, with discussion and analysis, your company believes a consensus on a uniform approach could be reached?
3. Which areas do you believe that your company's uniqueness will make it difficult to reach agreement on a common approach?

Based on the survey responses and the discussions at the workshops, this "Uniformity Report" has been prepared for discussion at the December 4, 2015 S-MAP Workshop #4.

⁴ Demographic or system differences refers to differences in size, organization structure, commodities, customer base, electric and gas system configurations, number of employees, geography of service territory, weather affecting the territory, urban vs suburban vs rural service territory, etc.

2.4. Report Structure

The Uniformity Report identifies the following:

- **The utilities' risk management frameworks⁵**
A discussion of the common approaches to addressing public and employee safety and security.
- **Areas of commonality**
Areas where a uniform approach can be adopted or worked out.
- **Areas of uniqueness**
Areas where each company needs to retain its approach due to differences in demographics and systems or where the costs of uniformity outweigh the benefits.
- **Areas for future consideration**
Areas where there will continue to be discussions to better understand the potential for uniformity
- **Commitment**
The utilities' agreed to continue to collaborate and share knowledge to enhance risk management practices and to look for additional areas of commonality as the understanding of risk management and the associated tools and processes evolve.

⁵ Each utility uses a framework to manage risks. The frameworks incorporate, among other things, the risks being addressed (e.g. Safety, Financial), risk categories (e.g. wildfire, gas explosion), matrices, models, and algorithms.

3. The Utilities' Risk Management Frameworks

The overall objective of the utilities was to identify common practices and approaches to assess, manage and mitigate **public and employee safety and security** risks. From the work in the initial utility uniformity workshop it became clear that the overall risk frameworks employed by each company are substantially similar. Each framework addresses: safety and other risks, risk types (e.g., wildfire, cyber, gas explosion) affecting each utility, and actions taken to mitigate the risk types. There was also recognition that there are differences in how the frameworks are implemented and, regardless of the differences, the same objectives are being achieved.

Throughout the General Rate Case (GRC) process, each company describes how it uses its risk management framework, practices and approaches to address types of risk (**Risk Categories**). The Risk Categories include:

- Wildfire
- Gas Explosion
- Work site accidents/employee, public, and contractor safety
- Accident involving contact with electrified asset
- Blackout/major customer outage
- Data breach (e.g., cybersecurity)

Even though there are differences in the size, financial capability, geography, commodities, structures and business processes of the utilities, each of the utilities has implemented processes and practices to address, where applicable, the risk categories noted above. Therefore, while there may be variances in approach, the end result, namely managing the right risks, is the same.

4. Areas of Commonality

As noted above, there were differences in how the specific tools and practices were used by the utilities to address risks. During the utility workshops each company's risk framework and the associated tools, methods and models used to support their respective risk frameworks were reviewed. Based on the discussions, the collaborative work of the companies, and the review of the risk frameworks, the participants were able to confirm where the adoption of a common standard was obvious and identify ones where, with dialogue, areas of uniformity could be adopted.

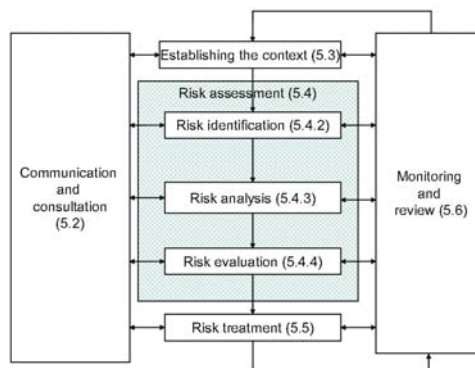
Obvious areas of commonality

1. ISO 31000 (Risk Management – See Figure 1) and 55000 (Asset Management)

As noted above, one of the CPUC objectives is to have common guidelines or standards across the utilities. Use of widely accepted international standards supports the Commission's objective and aligns with the utilities' current practices.

Each utility agreed to use as a foundational element on which to build its risk management framework, the **principles and guidelines** (tenets) of ISO 31000 and ISO 55000.

Figure 1 – ISO 31000 Framework



() indicates reference to ISO 31000

SED's current focus is on public and employee safety and security risks. ISO 55000, among other things, addresses a wider array of asset-related risks; therefore, there may be a need to clarify how broad the foundational use of ISO 55000 should be in order to balance with the Commission and SED's focus on

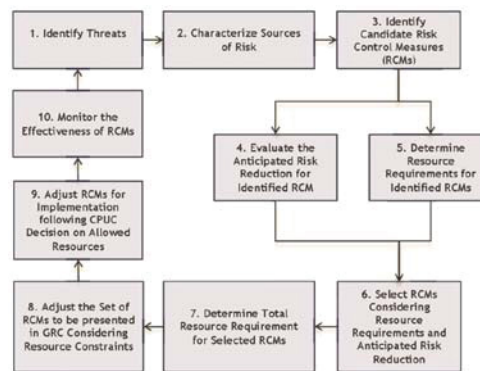
public and employee safety and security.

Areas where utility workshop discussions led to additional agreement on uniformity

2. Cyclo Process

The Cyclo 10-step process (See Figure 2) was presented in PG&E's 2014 GRC. The utilities incorporated the 10 steps into their risk frameworks. Each utility agreed to continue to support the use of the Cyclo 10-step process. The utilities noted that while all of their processes use the 10-step Cyclo process their graphical representation of the process will vary by company.

Figure 2 Cyclo 10-Step Process



3. 7x7 Matrix

Items 3 through 10 below address elements of the matrices used by the companies to address their risks. Each element of the risk matrix was discussed in detail to identify areas of commonality.

Over the last few years, the utilities have each tested various matrices, including 3x3, 5x5 and 7x7, to categorize risks. The 3x3 and 5x5 matrices were not sufficiently granular to adequately differentiate the impact and likelihood of the risks. As a result, the utilities agreed to continue to use the 7x7 matrix.

4. 7x7 Impact Level Descriptors

There were slight differences in the descriptions of the impact levels within the 7x7 matrices. Through discussions, the utilities agreed to standardize these descriptors in an effort to facilitate communications of the impacts of the various risks they manage.

1	2	3	4	5	6	7
Negligible	Minor	Moderate	Major	Extensive	Severe	Catastrophic

5. 7x7 Impact Categories

Based on individual differences, priorities and requirements, slightly different categories have been used for evaluating impacts within the 7x7 matrices. Despite those differences, all the utilities include impact categories pertinent to the risks the various GRC parties are concerned with (e.g., safety, reliability and environmental). Below are each of the utility's impact categories as presented in their S-MAP testimonies:

PG&E: Safety, Environmental, Compliance, Reliability, Trust, and Financial.

SCE: Safety, Environmental, Compliance, Reliability, and Financial.

SoCalGas and SDG&E: Health, Safety and Environmental; Regulatory, Legal and Compliance; Operational and Reliability; and Financial.

6. 7x7 Likelihood Level Descriptors

Although the utilities had slight differences in describing the likelihood levels in their 7x7 matrices, through discussions, agreement was reached to standardize these descriptors in an effort to facilitate communications of the likelihood for the risks they manage.

1	2	3	4	5	6	7
Remote	Rare	Infrequent	Occasional	Frequent	Regular	Common

7. 7x7 Likelihood Criteria

Agreement was also reached to standardize the use of the same likelihood criteria for the various levels in their 7x7 matrices as follows:

1	2	3	4	5	6	7
Remote	Rare	Infrequent	Occasional	Frequent	Regular	Common
Once every 100+ years	Once every 30-100 years	Once every 10-30 years	Once every 3-10 years	Once every 1-3 years	1-10 times per year	>10 times per year

8. 7x7 Impact Criteria

The criteria within each impact category cell reflects the individual demographics and systems of each of the utilities, which vary due to natural differences in operating environments, service territories and other factors. As such, the utilities will need to maintain flexibility to differentiate the levels of impact as deemed appropriate by the individual utility. For example, PG&E's criteria for an extensive reliability impact varies from SoCalGas and SDG&E's as seen below:

PG&E	
Extensive (5)	Location: Impacts multiple critical locations or customers; or Duration: Disruption of service greater than 10 days; or Customer Impact: Unplanned outage (net of replacement) impacts more than 10k customers; or EO: 100 thousand total customer hours, or more than 10 thousand MWh total load; GO: 100 thousand total customer hours, or reduction of capacity greater
SoCalGas/SDG&E	
Extensive (5)	> 50 K customers affected or impacts multiple critical locations or customers; substantial disruption of service greater 10 days.

Therefore, there was agreement that the impact criteria within each cell (1-7) of the 7x7 matrix should be defined by each utility.

However, there was agreement that each company shares the same safety concerns and, therefore, there was agreement to adopt common safety impact criteria as follows:

Impact Level	Description
Catastrophic (7)	Fatalities: Many fatalities and life threatening injuries to the public or employees.
Severe (6)	Fatalities: Few fatalities and life threatening injuries to the public or employees.
Extensive (5)	Permanent/Serious Injuries or Illnesses: Many serious injuries or illnesses to the public or employees.
Major (4)	Permanent/Serious Injuries or Illnesses: Few serious injuries or illnesses to the public or employees.
Moderate (3)	Minor Injuries or illnesses: Minor injuries or illnesses to many public members or employees.
Minor (2)	Minor Injuries or illnesses: Minor injuries or illnesses to few public members or employees.
Negligible (1)	No injury or illness or up to an un-reported negligible injury.

9. 7x7 Matrix Impact vs Likelihood Absolute vs Continuous Values

Each of the utilities has adopted a different approach as to whether the values (1-7) within the cells of the matrix are either discrete or continuous.

	Discrete	Continuous
Impact	SCE, SoCalGas/SDG&E	PG&E
Frequency/Likelihood	SoCalGas/SDG&E	PG&E, SCE

The ultimate objective is to use continuous values, but the data available today and the quality of the data is not conducive to using continuous

values. Therefore, there was agreement to maintain flexibility in applying the 7x7 evaluations based on the quantity and quality of data available.

5. Areas of Uniqueness

As noted above, the utilities held an additional workshop and undertook further analysis of their scoring algorithms. The utilities' approaches used to score risks (i.e., algorithms) and the tools and methods used to address particular Risk Categories (e.g., Wildfires) were reviewed and tested. In each case (scoring risks, and tools and methods for Risk Categories), individual demographics and systems made it difficult to use a common standard. Each of these areas is discussed further below.

- Risk Scoring Algorithm:** There are many areas where the combined utilities are regulated, yet each use slightly different approaches to compliance. For example, today each utility uses different measures of reliability (e.g., SAIFI, SAIDI, CAIDI) even though all the companies are focused on providing highly reliable service. Similarly, within energy procurement there is agreement that there should be a 95% confidence that each company's electric commodity rates will not change by plus or minus 10%, but different companies use different models to calculate their confidence. In the case of risk scoring, each utility uses its own algorithm to determine where a particular risk will fall in its risk matrix. In much the same manner as the above examples, the different scoring approaches produce similar results. Once top risks are identified, each utility determines the strength of current controls in mitigating the risk (i.e., the current residual risk) and identifies additional mitigations if needed. Due to the fact that many of these assessments are subjective by nature, a proper calibration across the organization is important, and the ranking of the risks reflects differences in calibration. As a result, the utilities concluded that for the foreseeable future, each utility will continue to use its own algorithm.
- Tools and methods to score Risk Categories:** Each utility has developed tools and methods that are unique to how particular risks are affected by their demographics and systems. For example, while all utilities have Wildfires as a risk category, how the level of risk and the resulting mitigations are developed will vary by utility because of differences in topography, weather, system infrastructure, etc. To further illustrate the point, wind is a key driver of wildfire risk. In Southern California, Santa Ana winds increase wildfire risk, but Northern California is not subject to this weather condition. In Northern California, high winds typically occur in winter and are often accompanied by rain, presenting less of a wildfire risk. Therefore, the utilities believe that allowing flexibility in what tools and methods are used to assess and manage specific risk category mitigants will be more effective than using one approach for all utilities and all risk categories.

6. Areas for Future Consideration

The utilities found the discussions and reviews productive and believe that, as practices mature, there will be other areas where common or similar approaches can be adopted. Three specific areas considered for future discussion are risk tolerance, risk reduction benefit per dollar invested and risk taxonomy. The utilities recognize that while addressing each of these three areas is complicated and will take time, the benefits of continuing to discuss and share approaches in each area will make some of the areas of uniqueness (e.g., algorithms) less significant. Each of these areas is discussed further below:

- **Risk Tolerance:** The utilities discussed the need for a risk tolerance framework where each utility could develop its own risk tolerance to allow the CPUC and intervenors to develop their own perspective on each utility's proposed risk tolerance. The utilities believe that a risk tolerance determination will have to allow for demographic and system differences and that the potential exists for some common approaches. The utilities also believe that further thinking on risk tolerance will emerge from the risk tolerance discussions and efforts.
- **Risk reduction benefit per dollar invested:** Each of the utilities has risk-informed practices and processes in place with which to fund risk mitigants. Additionally, the utilities have developed approaches that fit with their financial structures. While none of the implemented funding methods is currently capable of generating a risk reduction benefit per dollar invested, all of the utilities agreed that with further discussions, a common and transparent approach for evaluating risk mitigation effectiveness can be established utilizing the existing 7x7 matrix. This potential approach would evaluate various mitigation strategies at the risk level to demonstrate how a collection of controls and mitigations contribute to reduction of a specific risk. Additionally, the utilities agreed that cross-risk prioritization of mitigations is more challenging and will require more efforts to establish a methodology that can optimize spend based on risk reduction benefits across an entire organization. The utilities also believe that while these efforts may result in some common approaches, differences in demographics and systems will have to be considered.
- **Risk Taxonomy⁶:** Today, SDG&E and SoCalGas have begun implementing a common risk taxonomy across all of its operations and risks. This is an area where all the utilities see a potential for achieving standardization in the classification of risks and operations, but further dialogue is required.

⁶ A risk taxonomy is a comprehensive, common and stable set of risk categories that is used within an organization.

7. Commitment

To conclude, the utilities believe that because of the focus on safety and security risks, the utilities' ability to reach consensus on key elements that pertain to such risks supports the Commission's objective of standardization. In particular, the utilities' agreements on the consistent application of the 7x7 matrix to evaluate the likelihood and safety impacts of any given risk will provide the Commission and parties with common elements to facilitate evaluations and discussions of the various risks that the utilities manage through the GRC process.

While the common elements will facilitate the ratesetting process, there are important future opportunities that will allow more complex issues of risk tolerance and risk quantification to be addressed through dialogue and shared learning.

Furthermore, the utilities found the uniformity discussions valuable and have agreed to continue to hold sessions in order to explore and share innovations and to be able to continue to develop common standards and guidelines that support movement towards implementing leading practices.

(END OF ATTACHMENT 2)